Title
Essays on inference and strategic modeling

Author
Gillen, Benjamin J.

Publication Date
2010

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Essays on Inference and Strategic Modeling

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Benjamin Joseph Gillen

Committee in charge:

Professor Allan Timmermann, Chair
Professor Vincent Crawford
Professor Graham Elliott
Professor Harry Markowitz
Professor Joel Sobel
Professor Rossen Valkanov

2010

The dissertation of Benjamin Joseph Gillen is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2010

DEDICATION

To Sharin, for filling my life with sweetness.

To my parents for all their support and guidance through the years.
To Dad, who gave me my first econometrics problem when I was eleven.
To Mom, for making sure I didn't take it too seriously.

To La Jolla Shores, Mt. Woodson, San Elijo, San Onofre, Scripps Pier,
Kassia Meador, Dennis Murphy, Donald Takayama, Joel Tudor,
and the good people at the Karl Strauss Brewing Company
for helping clear my mind.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I owe the greatest thanks to my adviser Allan Timmermann for his guidance and support throughout my graduate experience. As a mentor and teacher, he is a model for economists and human beings alike and especially for economists who are human beings.

I'm also especially grateful to Vince Crawford, who has been immensely supportive throughout my graduate career. In addition to being the nicest economist I've met, his insights and encouragement have been invaluable in developing my work.

In general, I received a lot of help from the theorists here at UCSD. Joel Sobel has been instrumental in developing my theoretical skills. David Miller provided a receptive sounding board and directed me toward work in auctions. Nageeb Ali has also been very helpful in developing ideas and recommending new directions. I blame Chapter 3 on Navin Kartik, who introduced me to Doug Bernheim's work on conformity.

Since Allan wouldn't let me out with a pure theory dissertation, I also required a great deal of help from the econometricians. I'm grateful to Graham Elliott for consistently asking hard questions of great worth. Ivana Komunjer's expertise in identification arguments was extremely helpful. Andres Santos was instrumental in developing my auction estimation strategy. While technically a graduate student, Hiroaki Kaido's help warrants recognition with the faculty, as I'd not have been able to complete this work without it.

When I started work on my dissertation, I expected to write entirely in the field of financial economics rather than just one chapter. I benefited greatly in developing the chapter on covariance matrix estimation from Rossen Valkanov's comments and encouragement. I'm also particularly grateful for the opportunity to work with Harry Markowitz, whose esteemed reputation is outstripped only by his true quality.

From studying for the quals to reviewing job market papers and practicing our talks, I've benefited greatly from having good friends and colleagues with me every step of the way, especially Bryan, Choon, Jake, and Mike.

VITA

1999	B. A. in Economics, Yale University
2005-2010	Teaching Assistant, University of California, San Diego
2010	Ph. D. in Economics, University of California, San Diego

PUBLICATIONS

Benjamin J. Gillen and Harry M. Markowitz, "A Taxonomy of Utility Functions", Variations in Economic Analysis: Essays in Honor of Eli Schwartz, 2009.

FIELDS OF STUDY

Major Field: Economics

Studies in Behavioral Economics and Game Theory:
Professors Vincent Crawford and Joel Sobel

Studies in Econometrics:
Professors Graham Elliott and Allan Timmermann

Studies in Financial Economics:
Professors Harry Markowitz, Allan Timmermann, and Rossen Valkanov

ABSTRACT OF THE DISSERTATION

Essays on Inference and Strategic Modeling

by

Benjamin Joseph Gillen

Doctor of Philosophy in Economics

University of California, San Diego, 2010

Professor Allan Timmermann, Chair

This dissertation presents three stand-alone contributions to econometric inference and the analysis of strategic behavior.

Chapter 1 develops a structural econometric framework for first-price auctions that generalizes the assumption of Bayesian Nash Equilibrium within the context of a level-k behavioral model. The level-k model nests equilibrium by allowing bidders to best respond to heterogeneous beliefs about opponents' bidding strategies. I characterize conditions for identification of the distribution over valuations and bidder-types in populations with heterogeneous behavioral strategies. I propose a semi-nonparametric maximum likelihood estimator, establishing nonparametric consistency with an upper-semicontinuous population likelihood function, which I compute using a generalized expectation maximization algorithm.

Presenting evidence from a pilot study of vintage computer auctions, I find a high level of bidder sophistication in the field. I also characterize expected revenues in first price auctions with level-k bidders, establishing a partial identification result for expected revenues in unidentified behavioral models. Empirical evidence suggests a misspecified equilibrium optimal reserve price could reduce expected revenues 30% relative to an unbinding reserve price.

Chapter 2 introduces new Bayesian methods adapted to estimating a large-dimensional covariance matrix. I analyze the return generating process using an unrestricted factor model of covariance, imposing structure on the covariance matrix through prior beliefs on the parameters governing this process. I use these methods to provide an empirical Bayesian foundation for a general class of shrinkage estimators and use the shrinkage interpretation to characterize prior beliefs that optimize a posterior objective function. This estimation strategy delivers lower finite-sample loss than existing estimators in Monte Carlo simulations and performs well in minimum variance portfolio selection exercises.

Chapter 3 analyzes conformist tendencies for a population in which individuals gain utility by mimicking the average behavior, characterizing norms by the mean behavior, thus introducing an endogenous mechanism for establishing social norms. Under this specification, social preferences generally give rise to more concentrated behavior and a conformist pool forms when social preferences are sufficiently prominent. In addition to illustrating the determinants of conformist behavior with an endogenous reference point, these findings support applied work inferring social norms from average behavior.

Chapter 1

Identification and Estimation of Level-k Auctions

I develop a structural econometric framework for first-price auctions by generalizing the assumption of Bayesian Nash Equilibrium within the context of a level-k behavioral model, which nests equilibrium by allowing bidders to hold heterogeneous beliefs about opponents' bidding strategies. While behavioral heterogeneity causes identification to fail under benchmark equilibrium conditions, independence and exclusion restrictions recover identification of the joint distribution over valuations and bidder-types in heterogeneous populations. Establishing consistent maximum likelihood sieve estimation with an upper semicontinuous population log-likelihood function leads to a natural semi-nonparametric maximum likelihood estimator based on Legendre polynomials. The level-k model introduces a mixture structure to the estimation problem, requiring a generalized expectation maximization algorithm. Presenting evidence from a pilot study of vintage computer auctions, I find a high level of bidder sophistication in the field. To further apply the econometric framework, I characterize expected revenues in first price auctions with level-k bidders, establishing a partial identification result for expected revenues in unidentified models. An empirical analysis of USFS timber auctions finds that a misspecified equilibrium optimal reserve price could reduce expected revenues up to 30% relative to an unbinding reserve price.

1.1 Introduction

Empirical studies of auctions focus on estimating the distribution over valuations held by a representative bidder for the item being sold at auction. This distribution plays a key role in counterfactual analysis and in characterizing the effect of policy (for example, the auction's reserve price) on the expected revenue generated by the sale. In the estimation problem, an econometrician interprets data on bidders' characteristics, the object for sale, and the bids themselves, using a structural econometric model to infer the population distribution over latent values. The model links an individual bidder's unobserved valuation to their observed bid by imposing structure on the dependence in valuations across bidders, individual risk preferences, and the strategic beliefs bidders hold when choosing their optimal bid. A substantial literature leverages the Bayesian Nash Equilibrium (BNE) model of behavior to simultaneously impose structure on behavior and beliefs that allows for the unique identification and optimal estimation of the distribution of latent bidder valuations from the observed distribution over bids. Seminal contributions driving this line of research are due to Laffont and Vuong (1993), Donald and Paarsch (1996), Guerre, Perrigne, and Vuong (2000), and Athey and Haile (2002), with broad surveys presented in Athey and Haile (2005) and in Paarsch and Hong (2006)'s textbook.

In this paper, I generalize the empirical analysis of auctions by nesting Bayesian Nash Equilibrium behavior within the level-k behavioral model proposed by Crawford and Iriberri (2007a) based on the theoretical development of level-k cognitive hierarchy models developed by Costa-Gomes, Crawford, and Broseta (2001) and Camerer, Ho, and Chong (2004). The level-k model retains the rationality assumption that players best respond to beliefs about opponents' strategies but allows these beliefs to be heterogeneous and drawn from a structured hierarchy that gives rise to a mixture of behavioral types in the population. Beyond rationalizability, the level-k model restricts a player's bidding strategy to a set of decision rules, or types, defined through an iterated belief hierarchy anchored in an uninformative Level-0 (L0) model of opponents' behavior. The Level-1 (L1) player-type bids optimally based on the belief that all their opponents follow the

L0 strategy, the Level-2 (L2) player best responds to the belief that all their opponents follow the L1 strategy, and so on. By nesting the BNE behavioral model within its hierarchy, the more general level-k framework provides a robust foundation for inference on the distribution over valuations. Further, the generalization facilitates studying the distribution over bidder-types and characterizes the level of strategic sophistication in the population using field data. Incorporating the additional dimension to the model allows equilibrium to be tested against the directed alternative of a level-k behavioral model, providing external validation of the evidence accumulated in lab settings.

When Bayesian Nash Equilibrium is augmented with a level-k model of behavior, existing identification arguments no longer apply and identification generally fails in benchmark settings. In the BNE behavioral model, all individuals best respond to beliefs consistent with the observable distribution over bids, providing a key to the BNE identification argument that is not available in the level-k identification argument. The identification analysis begins by considering the setting in which the econometrician has sufficient information to identify the population distribution over bids in a homogenous population of known bidder-types. This setting highlights the link between bids and valuations conditional on the behavioral type, while controlling for potential issues related to identifying the distribution over bidder-types. Having established identification in a homogeneous population of bidders, identification fails in heterogeneous populations due to the need to identify both the distribution over valuations and the distribution over behavioral types. This expanded model has a dimensionality that exceeds the dimensionality of the information set, resulting in an incompletely identified model under the benchmark specification where the econometrician observes only the distribution over bids. As a consequence, for any distribution over bidder-types, there exists a distribution over latent valuations consistent with the observed distribution over bids.

While the seminal identification argument in Guerre, Perrigne, and Vuong (2000) leads to the natural derivation of an optimal indirect estimator, my identification results do not yield a clear estimation strategy for the level-k auction model outside of trivial settings. To address estimation of the model, I propose a consistent Semi-Nonparametric Maximum Likelihood (SNP-ML) estimator based on the Legendre polynomial sieve proposed by Bierens (2006) and adapted to a Simulated Nonlinear Least Squares (SNLS) estimator for equilibrium auction models by Bierens and Song (2007). The mixture of bidder-types in level-k auctions gives rise to an upper semicontinuous likelihood criterion function, requiring a more general uniform strong law of large numbers than applied by Bierens (2006) and by Bierens and Song (2007). I adopt a specialized version of the uniform strong law of large numbers developed in Artstein and Wets (1995), which is closely related to Hess (1996)'s results, using techniques based on epiconvergence that have proven particularly useful in the nascent study of set estimators for partially identified models.[1] Consistency of the SNP-ML estimator with upper semicontinuous likelihood functions also allows the econometrician to flexibly control for observed auction heterogeneity in the model, formally extending Donald and Paarsch (1996)'s maximum likelihood consistency argument to infinite dimensional semi-nonparametric auction problems. As Donald and Paarsch (1996) show, this flexibility comes at the expense of non-standard asymptotic analysis that hinders characterizing the distribution of test statistics based on the SNP-ML estimator. In computation, the heterogeneous behavioral model results in a mixture structure for the log likelihood criterion function so that a direct likelihood maximization exercise is computationally infeasible. To address this complication, a generalized expectation maximization algorithm partitions parameter space into a set of parameters for which closed-form solutions are readily available and another set requiring numerical methods for optimization. Partially maximizing over those parameters for which closed-form solutions are available before maximizing over the full set of parameters greatly reduces computation time, rendering calculation of the estimator feasible.

[1]The analysis in Kaido (2009) uses weak epiconvergence arguments in deriving a set of distributional results that unify set estimation techniques based on vector support functions developed by Beresteanu and Molinari (2008) and approximate minimizers to criterion functions developed by Chernozhukov, Hong, and Tamer (2007). Other recent works in econometrics utilizing weak epiconvergence to develop point estimation results in non-standard problems include Chernozhukov and Hong (2004) and Han and Phillips (2006).

In an application of the econometric model, I consider the mechanism design problem of selecting the reserve price to maximize expected revenues in a level-k model. This analysis extends the level-k bidding model for auctions with a reserve price developed in Crawford, Kugler, Neeman, and Pauzner (2009) to a general distribution over valuations with more than two bidders participating. I also address the cases where the mechanism designer does not know the composition of the bidding population and, even worse, when the distribution over bidder-types is unidentified. Solving for the expected revenue to the seller at a fixed reserve price, I present first order conditions for the optimal reserve price similar to those developed by Myerson (1981). In the setting where the distribution over bidder-types is not identified, the seller's expected revenues are only partially identified. As such, identification of the optimal reserve price is similar to the result in Haile and Tamer (2003)'s study of jump bidding in English auctions. Specifically, the optimal reserve price under the unidentified behavioral model belongs to a closed, bounded set. To close the mechanism designer's decision problem and to prescribe a unique, robust optimal reserve price, I introduce min-max preferences over the ambiguity generated by the lack of a fully identified model. These preferences adapt naturally to the level-k model as the minimum conditional expected revenue can be calculated by only considering the expected revenues in a homogenous population of bidder-types.

I estimate the model using data from USFS timber auctions, with empirical results that contradict the equilibrium stylized fact that expected revenues in US Forestry Service timber auctions could be increased by introducing a binding reserve price. In particular, given a sufficiently large population of unsophisticated bidders, expected revenues are maximized by setting the reserve price equal to the government's appraised value for the tract. This result rationalizes the current USFS policy of implementing non-binding reserve prices without having to appeal to non-revenue motives.

I then turn to an empirical study of bidder sophistication in the field. Given the robust experimental analysis of bidder behavior in the lab, the goal of this application is to provide evidence testing the external validity of these results. To this ends, I present results from a pilot study of bidding behavior observed in a series of sealed-bid, first-price, auctions for vintage computer equipment from the Alameda County Computer Resource Center. While estimated results from this sample support the equilibrium bidding model, additional data points are required to generate more robust findings. As such, the sample size is quite limited and so the findings are best considered the pilot for a broader field study on auction behavior.

After a brief review of related literature in the next section, section 1.3 introduces the auction model and identification problem, illustrating the effect of behavioral misspecification on inference in a simple parametric example. Section 1.4 addresses the question of identification in homogeneous populations, stepping through the proposed hierarchy of behavioral types to progressively develop the intuition underlying the general identification argument. I extend these results to heterogeneous populations in section 1.5, beginning with the setting where each individual's type is observed by the econometrician and exploring exclusion restrictions that identify the distribution over types. Section 1.6 discusses estimation of the model in a semi-nonparametric framework, presenting the data generating process for bids and establishing consistency for the estimator. Section 1.7 addresses the mechanism design problem, with section 1.8 presenting estimates for optimal reserve pricing from USFS timber auctions. Section 1.9 discusses the pilot field study of bidding behavior before concluding.

1.2 Related Literature

The level-k model's mixture-of-types framework is rooted in earlier work by a number of researchers, including Stahl and Wilson (1995), Nagel (1995), and El Gamal and Grether (1995) with additional theoretical development by Costa-Gomes, Crawford, and Broseta (2001) and Camerer, Ho, and Chong (2004). Ho, Camerer, and Weigelt (1998) and Bosch-Domenech, Montalvo, Nagel, and Satorra (2002) apply the model to analyzing behavior in beauty contest games, presenting some of the first evidence that players rarely reach beyond the second level of the strategic hierarchy. Crawford (2003), Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007b) explore a number of applications for the model relating to information transmission. Ivanov, Levin, and Niederle (2008) present experimental evidence using common value auctions that questions whether individuals best respond to beliefs, illustrating the difficulty in advancing individuals to higher levels of the hierarchy.

The problem I consider is closest to an application considered by Aradillas-Lopez and Tamer (2008), who present general results for identification in strategic models under level-k rationalizability, as defined by Bernheim (1984) and Pearce (1984). Aguirreagabiria and Magesan (2009) extend Aradillas-Lopez and Tamer (2008) by evaluating identification when agents best respond to non-equilibrium beliefs about their opponents' strategies in dynamic games. Aradillas-Lopez & Tamer's set identification result provides an upper bound on the cumulative distribution function for the distribution over valuations (since no one would rationally choose to bid more than they thought the item was worth), but they go on to show that the identified set also includes any sufficiently regular distribution that first order stochastically dominates this bound. The additional structure of Crawford and Iriberri (2007a)'s level-k auction model provides significant identification restrictions beyond those of rationalizability by placing an upper bound on an individual's bid shade and, consequently, a lower bound for the cumulative density function over valuations. By exploiting this additional structure, the level-k model generally provides a much tighter identified set than is available in Aradillas-Lopez & Tamer's model, and admits point-identification for both the distribution over valuations and the distribution over bidder types under viable information specifications.

This paper contributes to the growing literature analyzing the econometrics of strategic models with non-equilibrium behavior. Among the first statistical models of non-eqilibrium behavior, the Quantal Response Equilibrium (QRE) due to McKelvey and Palfrey (1995, 1998) provides a mechanism for incorporating statistical noise into individual behavior at games. Haile, Hortacsu, and Kosenok (2008) and Goeree, Holt, and Palfrey (2005) illustrate the need for structure on that noise to obtain empirical restrictions, but by exploting the model's generality, Rogers, Palfrey, and Camerer (2009) introduce a heterogeneous model of QRE with a structured error term that nests cognitive hierarchy behavior. Goeree, Holt, and Palfrey (2002) solve the QRE for auction models, illustrating QRE's ability to generate overbidding in the presence of asymmetric distribution over valuations. Bajari and Hortacsu (2005) develop a structural econometric model for interpreting experimental auction data based on the QRE solution and also introduce an alterative non-equilibrium approach based on modeling behavior as an adaptive learning strategy.[2] In analyzing behavior at English auctions with jump bidding, Haile and Tamer (2003) develop a model that only imposes individual behavior be rationalizable, which results in a partially identified distribution over valuations. Haile and Tamer (2003) derive tight bounds for this identified set and present empirical estimates for the set-identified optimal reserve price in US Forestry Service timber auctions.

The identification results for auctions with an unknown distribution over non-equilibrium strategic beliefs are analogous to the results when equilibrium behavior is subject to unknown risk aversion. In the equilibrium identification analysis, risk neutrality or risk aversion of known form provides a key identifying assumption. Campo, Guerre, Perrigne, and Vuong (2000) present the non-identification result for the benchmark equilibrium model under parametric HARA utility specification with an unknown risk aversion coefficient. In this case, it is possible to identify an observationally equivalent distribution over valuations from the observed distribution over bids for any value of the risk aversion coefficient, resulting ina paritally identified model. This partial identification result requires additional information to identify risk aversion and is closely related to the need for additional information to identify the distribution over types in a level-k auction model. As such, I follow the approach proposed in Lu (2004), Bajari and Hortacsu (2005), Perrigne and Vuong (2007), and Guerre, Perrigne, and Vuong (2009), who recover identification of the bidder's utility function and the distribution over val-

[2]While not the central focus of the analysis, Appendix 1.A.3 briefly discusses identification under the Goeree, Holt, and Palfrey (2002) QRE auction model, deriving similar identification conditions to those in the Level-k behavioral model.

uations by exploiting testable restrictions originally proposed by Athey and Haile (2002). Since strategic uncertainty changes with the level of competition, testable restrictions generated by exogenous variation in the number of bidders participating in each auction provide information about an individual's joint risk preferences and strategic beliefs.

Maximum likelihood estimation of auction models is well known to present a computational burden. These challenges have motivated researchers to explore alternative estimation techniques such as simulated nonlinear least squares, originally developed by Laffont, Ossard, and Vuong (1995) and extended by Bierens and Song (2007), and simulated method of moments, applied in studies by Laffont and Vuong (1993) and Li (2005). The simulated nonlinear least squares methods are not available for estimating the level-k auction model due to their reliance on revenue equivalence for equilibrium behavior across auction mechanisms, which Crawford, Kugler, Neeman, and Pauzner (2009) show fails in the level-k behavioral setting. Similarly, the moment conditions derived from individual rationality used in simulated method of moments estimation can't be immediately adapted to the mixture of behaviors observed in the level-k model.

1.3 The Auction Model and Identification

This section formally defines the auction model, characterizes the identification objective, and reviews the level-k behavioral model. In reviewing Guerre, Perrigne, and Vuong (2000)'s non-parametric identification result, I focus on bidders' strategic beliefs, proposing a slightly abbreviated proof technique. I then describe Crawford and Iriberri (2007a)'s level-k behavioral model in greater detail and state regularity conditions to ensure the level-k behavioral model results in continuous, monotonic bidding behavior. The section closes with an example illustrating the incorrect inference that results in a mis-specified behavioral model.

1.3.1 The Auction Model and Equilibrium Identification

I consider identification in the risk neutral symmetric Independent Private Values (IPV) specification of the general Milgrom and Weber (1982) first price auction model.[3] In the period t IPV auction, each player $i \in N_t \in \mathcal{N} \subset \mathbb{Z}$ observes N_t and the realization of a random variable, X_{it}, which has a commonly known distribution, $F_X(x)$, that is absolutly continuous over $[\underline{x}, \overline{x}] \subset \mathbb{R}^+$ with a strictly positive pdf $f_X(x)$. The variable $X_{i,t}$ specifies bidder i's latent valuation for the good being sold at auction t. Bidder i then chooses his bid, s_{it}, conditional on this valuation and pays the value of his bid in exchange for the good if he submits the highest bid in the auction. Given that f is continuously differentiable and bounded away from zero, Maskin and Riley (1984) show that there exists an equilibrium in strictly monotonic and continuously differentiable bidding strategies.

As is standard in auction identification problems, the econometrician's benchmark information set consists of the empirical distribution over bids, denoted $F_S(s)$. When the inverse bidding function exists, this distribution over bids is *generated* by the true distribution over valuations and the equilibrium bidding function, $\sigma_{Eqm,X}(x)$, when $F_S(s) = F_X\left(\sigma^{-1}_{Eqm,X}(s)\right)$. The econometric model under equilibrium then is defined entirely by the distribution over valuations, which is identified if it is the unique distribution that generates the observed distribution over bids. This definition of observational equivalence is standard in the nonparametric identification literature from Brown (1983) and Roehrig (1988)'s early

[3]In the second-price auction format with Independent Private Values, equilibrium is in weakly dominant strategies so that, for any belief of opponent's behavior, selecting a bid equal to the agent's valuation is a best response. Since the level-k bidding behavior is identical to the equilibrium weakly-dominant strategy in this setting, identification is inherited from existing results and the distribution over bidder-types is trivially non-identified. Maintaining the independence and private valuation assumption avoids known challenges to identification presented in Athey and Haile (2002) in Common Values and Affiliated Private Values models from benchmark analysis, leaving these problems for future work. Addressing the class of pure common values models that Fevrier (2008) studies under the equilibrium behavioral model would be particularly interesting given the relationship between level-k behavioral model in these settings and the notion of "cursedness" as developed by Eyster and Rabin (2005). The symmetry assumption is made largely for the purposes of tractability and can be relaxed in an extension of Brendstrup and Paarsch (2006). Common knowledge of the number of bidders in each auction provides a key to identification by evaluating the strategic trade-offs an individual makes when confronted with different levels of competition.

contributions to Manski (1995)'s Survey and the more recent work by Benkard and Berry (2006) and by Matzkin (2008). Given this definition, identification fails if there exists an alternative distribution over valuations, $F^* \neq F_X$, that is observationally equivalent to F_X, defined as follows.

Definition 1.1 (Equilibrium Observational Equivalence)
A structure, (F_X), coupled with the equilibrium bidding rule $\sigma_{Eqm,X}(x)$ is observationally equivalent to the structure (F_*) coupled with the corresponding equilibrium bidding rule $\sigma_{Eqm,*}(x)$ if:

$$F_X\left(\sigma^{-1}_{Eqm,X}(s)\right) = F_S(s) = F_*\left(\sigma^{-1}_{Eqm,*}(s)\right)$$

Now consider the identification argument in the equilibrium behavioral model. Dropping the t subscript unless needed for clarity, individual i's payoff is:

$$U_i(X_i, s_1, \ldots, s_N) = (X_i - s_i)\, 1_{\{s_i > \max_{j \neq i} s_j\}}$$

Conditional on X_i, the independence of valuations (and, consequently, of bids) implies player i's expected utility from the bid s_i is:

$$E\left[U_i(X_i, s_1, \ldots, s_N)\,|X_i\right] = (X_i - s_i)\, Pr\{s_i > \max_{j \neq i} s_j\} \tag{1.1}$$

Equilibrium analysis typically begins by hypothesizing a behavior for other players, solving for player i's best response to this behavior, and then finding a fixed point where everyone's behavior is consistent with rational beliefs. This analysis ensures that, first, all players are best responding to beliefs and, second, that those beliefs reflect the true joint distribution of behavior and valuations. While the first feature links the individuals' bids to bidder valuations, it is the second feature of equilibrium that links the econometrician's information set to the player's information set, providing the key to establishing identification. Given the empirical distribution over bids, symmetry and independence imply the true probability that player i will win the auction with a bid of s_i is given by the cumulative distribution of the highest bid among $N-1$ independent competing bidders: $Pr\{s_i > \max_{j \neq i} s_j\} = F_S(s_i)^{N-1}$. Since equilibrium requires player i's beliefs to match this empirical distribution over bids, the expected utility that player i seeks to maximize in equilibirum is:

$$E\left[U_i\left(X_i, s_1, \ldots, s_N\right) | X_i\right] = \left(X_i - s_i\right) F_S\left(s\right)^{N-1} \tag{1.2}$$

The first order conditions for optimal behavior establish identification of the IPV auction model by recovering the true valuation for any bid directly from the bid's value and the distribution over bids, F_S. These first order conditions give the inverse bidding function:

$$x_i = s_i + \frac{F_S\left(s_i\right)}{\left(N-1\right) f_S\left(s_i\right)} \tag{1.3}$$

This identification argument is more direct than the seminal argument in Guerre, Perrigne, and Vuong (2000), which analyzes the bid function transforming valuations into bids to illustrate the role of the Jacobian of the transformation in establishing identification. By focusing on beliefs, the proof highlights a key equilibrium feature: player i's expected utility depends only on their own independent private valuation and other player's actions, which are i.i.d. with a distribution contained in the econometrician's information set. Many existing results regarding the identification of auctions, notably those that incorporate parametric risk aversion (such as Campo, Guerre, Perrigne, and Vuong (2000) and Lu (2004)) as well as the asymmetric bidder model studied by Brendstrup and Paarsch (2006), can be similarly proved directly by exploiting this property. This more direct proof technique could be extended to state a set of sufficient conditions that can be applied to establish identification in a range of games with incomplete information, including principal-agent problems, coordination and search games.

The remainder of the paper assumes these following relevant properties of the IPV auction model:

Assumption 1.1 (IPV Auction Model) *Unless explicitly stated otherwise,*

a. $U_i\left(X_{it}, s_{1t}, \ldots, s_{Nt}\right) = \left(X_{it} - s_{it}\right) 1_{\{s_{it} > \max_{j \neq i} s_{jt}\}}$

b. *$X_{it} \sim_{iid} F_X$ which is absolutly continuous over $\mathcal{X} = [\underline{x}, \overline{x}] \subset \mathbb{R}^+$ with strictly positive pdf $f_X(x)$*

c. *N_t and F_X are common knowledge.*

1.3.2 The Level-k Behavioral Model

In Crawford and Iriberri (2007a)'s level-k auction model, a player draws their beliefs about other players' actions from a cognitive hierarchy and best responds to those beliefs, giving rise to a mixture of heterogeneous behavioral types in the population. The level-k behavioral model provides a parsimonious framework for modeling how people approach decision making in novel strategic by imposing structure on the set of possible beliefs governing individual behavior. Crawford & Iriberri consider two possible specifications for anchoring level zero (L0) beliefs and characterize the strategies for L1 and L2 players based on these anchoring beliefs. Players with higher levels of sophistication are not observed in the lab and, as such, are treated as effectively non-existent. A Random L0 player ($\mathrm{L0}_R$) bids uniformly over the set of possible valuations and a Truthful L0 player ($\mathrm{L0}_T$) truthfully reveals their expected valuation for the object on auction. These two anchoring beliefs each give rise to two belief hierarchies and behavioral types, with level-type $\mathrm{L}k_\tau$ best responding to the belief that everyone plays the $\mathrm{L}(k-1)_\tau$ strategy. When coupled with the equilibrium strategy behavior (corresponding to a level-∞ bidder type), this cognitive hierarchy allows for seven potentially different behavioral types with two of those types assumed to exist only in the imagination of other players. To summarize the behavioral assumptions:

Assumption 1.2 (Level-k Behavioral Model) *In the level-k behavioral model of auctions,*

a. *a player observes their valuation, the number of bidders participating in the auction, and is assigned to a bidder-type k,*

b. *the player of type k's strategy best responds to the belief that every other player plays according to the type $k-1$ strategy*

c. *the level-0 bidder-types' behavior is governed by an uninformative strategy: the $L0_R$ bidder-type bids uniformly over the set of valuations and the $L0_T$ bidder-type bids their valuation, and,*

d. *the level-∞ bidder-type bids in accordance with the BNE strategy.*

A similar approach to modeling strategic behavior is proposed by Camerer, Ho, and Chong (2004), who develop a cognitive hierarchy model for one-shot games rooted on an uninformative level-0 behavioral model. In contrast to Costa-Gomes, Crawford, and Broseta (2001)'s level-k behavioral model, players at the k-th level in the cognitive hierarchy do not believe that every other player follows the level $k-1$ strategy, but rather believes that there is a mixture of players following the level-0 through the level $k-1$ strategy. As such, while the k-th level behavioral type in the cognitive hierarchy is oblivious to the notion that people are playing at their own or higher levels of sophistication, they recognize heterogeneity among lower bidder-types and know the relative proportion of the lower-level bidding types. In a model that assumes no level-0 types exist in the population, the level-1 and level-2 cognitive hierarchy behavioral types are identical to the $\text{L}1_R$ and $\text{L}2_R$ bidder types defined above, with the only difference between the models realized at higher levels of sophistication. Experimental evidence for the IPV setting presented in Crawford and Iriberri (2007a)'s online appendix suggests the truthful hierarchy of types are not as prominent under the IPV model as in common values settings. For this reason, the application only considers estimating the hierarchy based on the Random Level-0 bidder-type. As such, the only applied difference in the two modeling approaches arises for the level-∞ bidder-type, who follows the BNE strategy in the level-k model but best responds to the empirical distribution over bids in the cognitive hierarchy model.[4]

[4]In their development, Camerer, Ho, & Chong introduce additional structure to the model by assuming the sophistication of a bidder as being drawn from a Poisson distribution, which captures the intuition of the cognitive hierarchy as arising from an iterative reasoning procedure. Given the identification results, this Poisson Cognitive Hierarchy model can be modeled as a restricted version of an unrestricted distribution over bidder-types. From a computational perspective, the level-∞ bidder-type under the cognitive hierarchy is easier to accommodate as their beliefs can be calculated directly from the empirical distribution over bids using a strategy similar to that proposed in Guerre, Perrigne, and Vuong (2000). Finding the fixed point where

1.3.3 Regularity Conditions

Regularity conditions on the distribution over valuations ensure the agents' bidding functions are strictly monotonic so as to preclude pooling behavior that would stymie identification and introduce atoms to the estimation problem. These conditions are similar to Myerson (1981)'s regularity condition that ensures the mechanism designer's revenue optimizing problem is well defined:

$$1 - \frac{d}{dx}\frac{1 - F_X(x)}{f_X(x)} > 0 \tag{1.4}$$

Myerson's regularity condition presents a restriction on the inverse hazard rate for the event that one bidder will have a valuation exceeding a reservation price. Ensuring the level-k bidding strategy is continuously differentiable and strictly monotonic requires a restriction on the inverse hazard rate for the event that the player would win the auction. These conditions strengthen regularity conditions in standard auction theory, with assumption L1.1.3 representing the level-k analog to the Myerson's regularity condition. For purposes of generality, the lemma places conditions on the level $(k-1)$ bidding strategy, though it is straight-forward to verify these conditions hold inductively as a condition on the primitive distribution over valuations.

Lemma 1.1 (Level-k Regularity Conditions) *Suppose:*

L1.1.1 *The level-$(k-1)$ bidding strategy is a twice continuously differentiable, strictly monotonic function of their valuations.*

L1.1.2 *F_X has k continuous, bounded derivatives over support $[\underline{x}, \overline{x}]$, and*

L1.1.3 *$\exists \xi > 0$ such that, for all possible valuations $x \in \mathcal{X}$,*

$$1 - \frac{d}{dx}\frac{F_X\left(\sigma_{k-1}^{-1}(x)\right)}{f_X\left(\sigma_{k-1}^{-1}(x)\right)\frac{d}{ds}\sigma_{k-1}^{-1}(s)\mid_{s=x}} > \xi \tag{1.5}$$

the beliefs for the equilibrium bidder-type are consistent with the empirical distribution over bids in the level-k model presents the central computational challenge in the empirical analysis.

Then the level-k bidding strategy, $\sigma_k(x)$ is a bounded, strictly monotonic, and continuously differentiable function of the bidder's valuation.

The proof of Lemma 1.1 consists of implicitly differentiating the first order conditions of the level-k bidder-type's optimization problem. The regularity condition in equation 1.5 ensures the derivative of the bidding function is continuous and bounded away from zero by ξ.

1.3.4 Example: Log-Normally Distributed Valuations

To illustrate the implications of the behavioral differences between the level-k and equilibrium auction models for inference, consider an auction with $N = 4$ players where valuations are log-normally distributed with mean parameter 0 and variance parameter 1. Each player first observes their valuations and is then independently assigned to one of five behavioral types: Random Level-1 ($L1_R$), Random Level-2 ($L2_R$), Truthful Level-1 ($L1_T$), Truthful Level-2 ($L2_T$), or Equilibrium (Eqm). Figure 1.1 plots the behavioral strategies for each of the behavioral types, all of which are strictly monotonic in the valuations and continuous over their range.

Figure 1.2 plots the distribution of bids for the different bidder types, illustrating that each of the types can be differentiated from one another using a series of stochastic dominance relationships inherited from the underlying bid functions. These relationships imply that, given the bidding distributions for every agent, the econometrician would have sufficient information to sort each agent into their respective bidder-type. In addition, figure 1.2 includes the unconditional bid distribution for a sample population where a bidder is assigned to bidder-types $L1_T$, $L2_T$, $L1_R$, $L2_R$, and Eqm with probabilities 60%, 5%, 15%, 5%, and 15%, respectively. The distribution over bids generated by this sample population corresponds to the econometrician's ex post information set from the auction in the benchmark informational setting.

What if the econometrician ignored the heterogeneous behavior, and instead estimated the distribution over valuations using a mis-specified equilibrium model of behavior? In Figure 1.2, the equilibrium bidder-type's distribution over bids has significantly thinner tails than the mixed-population distribution over bids, so the estimated distribution over valuations would have much fatter tails than the true distribution over valuations. Figure 1.3 illustrates the magnitude of this misspecification effect in the estimated distribution of valuations. The dashed line presents the true distribution of valuations, the solid green line presents the distribution of the winning bid, and the solid red line presents the estimated distribution of valuations under the BNE assumption. The bias in inference is severe: the true 95th quantile valuation is only 5 but the estimated 95th quantile is an order of magnitude larger, extending way off the boundary of the graph to nearly 50. This bias is driven by overbidding from lower-level types that, under the equilibrium model, can only be justified by exceedingly large valuations.

1.4 Identification in Homogeneous Populations

Characterizing non-parametric identification in the level-k auction model requires first establishing identification when all players behave homogeneously, following the strategy of a single bidder-type. This section analyzes each of the level-k bidder types separately, presenting the behavioral strategies and establishing conditions for identification to obtain in homogeneous populations. These results all for identification when the econometrician observes sufficient information to characterize each individual's distribution over bids, for example, by following the history of a given bidder in a large number of independent auctions. While this setting is unlikely to arise outside of a laboratory environment, it provides a basic set of results that applies in more realistic application where the econometrician can observe only the unconditional distribution over bids within the auction. As a preview of the results in this section, the following theorem, proved in the appendix, generally characterizes identification of a level-k bidder-type in homogeneous populations.

Theorem 1.1 (Identification in Homogeneous Populations)

Assume the conditions of Assumptions 1.1 and 1.2 and Lemma 1.1. Suppose the econometrician observes the distribution over bids, $F_{S_N,k}$, for a homogeneous *population of level-k bidders with N bidders participating in each auction, then the distribution F_X is uniquely identified.*

Suppose further that the econometrician observes the distribution over bids, $F_{S_{N^*},k}$, for a homogeneous population of level-k bidders with $N^* \neq N$ bidders participating in each auction, then the level-k behavioral model is testable through overidentifying restrictions.

The proof of theorem 1.1 follows from continuity of bidding behavior and non-negativity of the bid shade in the level-k behavioral model. The remainder of this section applies theorem 1.1 to the level-k bidding model proposed by Crawford and Iriberri (2007a), illustrating the intuition behind the identification argument while formally reviewing the behavioral model in progressively more complex settings.

1.4.1 Trivial Identification & Non-Identification Results

For several of the behavioral types, in particular the level-0 player-types, identification is either trivial or impossible. Whenever there is a single truthful bidder, that bidder's distribution over bids will stochastically dominate all other bidders' distributions over bids. Consequently, if the econometrician observes enough information to identify each individual bidder's distribution over bids, the distribution over valuations is trivially identified by the distribution over bids. On the other hand, if the population consists entirely of purely random bidders, whose behavior is independent of their latent valuations, then identification is obviously impossible regardless of the information obtained about the distribution of bids.

Lemma 1.2 (Level-0 Trivial Identification & Non-Identification)

A *Suppose there is at least one $L0_T$ bidder-type in the population and the econometrician observes the distribution over bids for each individual, then the distribution over valuations is identified.*

B *Suppose all bidders in the population are $L0_R$ bidder-types, then even if the econometrician observes the distribution of bids for each individual, the distribution over valuations is unidentified.*

In the example from section 1.3.4, the Random Level 1 ($L1_R$) bidder-type's strategy is linear in his valuation, presenting a setting where identification is only slightly more complicated than the truthful level zero bidder-type. Since his beliefs, and consequently his behavior, are invariant to changes in the distribution over valuations, the $L1_R$ bidder-type provides a simple context for formalizing the level-k behavioral model free of identification problems. Following Crawford and Iriberri (2007a) in adopting Krishna (2002)'s notation, denote the maximum bid submitted by players other than player i by the random variable Y_i. The $L1_R$ bidding function solves:

$$\sigma_{L1_R}(x) = \underset{\sigma: X \to S}{\operatorname{argmax}} \left(X_i - \sigma(X_i)\right) F_{Y_i}\left(\sigma(X_i)\right)$$

This maximization problem yields first order conditions given in Crawford and Iriberri (2007a)'s, Equation 14:

$$\left(X_i - \sigma(X_i)\right) f_{Y_i}\left(\sigma(X_i)\right) - F_{Y_i}\left(\sigma(X_i)\right) = 0 \tag{1.6}$$

The $L1_R$ bidder believes Y_i to be the maximum of $(N-1)$ uniformly distributed random variables, behaving as if Y_i has cdf and pdf $F_{Y_i}(s) = \frac{(s-\underline{x})^{N-1}}{(\overline{x}-\underline{x})^{N-1}}$ and $f_{Y_i}(s) = \frac{(N-1)(s-\underline{x})^{N-2}}{(\overline{x}-\underline{x})^{N-1}}$, respectively. These beliefs yield the best responding bidding function:

$$\sigma_{L1_R}(X_i) = \frac{N-1}{N} X_i - \frac{\underline{x}}{N}$$

Since the $L1_R$ bidding function is a linear transformation of the valuation, identification is immediate. Given the distribution over bids, with cdf denoted $F_{S,L1_R}$, the cdf and pdf for valuations are:

$$\begin{aligned} F_X(x) &= F_{S,L1_R}\left(\sigma_{L1_R}(x)\right) = F_{S,L1_R}\left(\frac{N-1}{N} X_i - \frac{\underline{x}}{N}\right) \\ f_X(x) &= f_{S,L1_R}\left(\sigma_{L1_R}(x)\right) \frac{N-1}{N} = f_{S,L1_R}\left(\frac{N-1}{N} X_i - \frac{\underline{x}}{N}\right) \frac{N-1}{N} \end{aligned} \tag{1.7}$$

Hence, a single-step estimation procedure can infer the distribution over valuations by shifting and schaling the estimated disttribution over bids.[5]

[5]While this result does not follow directly from Theorem 1.1, it is stated as a corollary since the $L1_R$ bidder type can be modeled as if he believes his opponents' bidding functions are given

Corollary 1.1 (Identification of $L1_R$ Bidder-Type Valuations)
Suppose that an econometrician observes the distribution over bids, $F_{S,L1_R}$, from a homogeneous population of $L1_R$ bidder-types, then the distribution over valuations, F_X, is identified.

1.4.2 Identification of Higher-Order Bidder-Types

The higher-order bidder-types best respond to beliefs that their opponents take into account the distribution over valuations in choosing their bid. These bidder-types follow more sophisticated strategies than described in the previous section, with the distinction becoming particularly salient when reserve prices are considered in section 1.7. For notational convenience, now add the assumption that the lower support of the distribution over valuations is zero.

Truthful Level 1 ($L1_T$)

The Truthful Level 1 ($L1_T$) bidder-type best responds to the belief that other players submit bids exactly equal to their valuations. As such, in equation 1.6, the $L1_T$ bidder behaves as if Y_i is the maximum of $(N-1)$ random variables drawn from the distribution for valuations, with cdf and pdf $F_{Y_i}(s) = F_X(s)^{N-1}$, and, $f_{Y_i}(s) = \frac{1}{N-1} F_X(s)^{N-2} f_X(s)$, respectively. Crawford and Iriberri (2007a) show these beliefs yield first order conditions:

$$X_i = \sigma_{L1_T}(X_i) + \frac{F_X(\sigma_{L1_T}(X_i))}{(N-1) f_X(\sigma_{L1_T}(X_i))} \tag{1.8}$$

Here, the conditions on F_X in Theorem 1.1 ensure that the implicitly-defined bidding function, $\sigma_{L1_T}(X_i)$, is well-defined, uniformly continuous, and strictly increasing in X_i.

Identification in this setting presents the first challenging result in the paper. The inverse bidding function in equation 1.8 characterizes the identified set

by: $\sigma_{L0_R}(x) = \overline{x} * F_X^{-1}(x)$. Note that, while applying kernel methods to the nonparametric estimation strategy would achieve the optimal rate of convergence in a homogeneous population, the technique cannot be transferred to the settings with more complicated bidding behavior.

as containing any distribution over valuations consistent with the observed distribution over bids. However, the inverse bidding function itself depends on the true F_X. For any distribution over valuations, F^*, subject to the regularity conditions in Theorem 1.1, define:

$$\sigma_{*,L1_T}^{-1}(s) = s + \frac{F^*(s)}{(N-1)\, f^*(s)}$$

The observational equivalence of F_X and F_* under the $L1_T$ behavioral model then requires:

$$F_X\left(\sigma_{L1_T}^{-1}(s)\right) = F_s(s) = F^*\left(\sigma_{*,L1_T}^{-1}(s)\right) \tag{1.9}$$

The identification argument establishes that any distribution F^* satisfying this relationship must be identical to F_X almost everywhere through a pair of contradictions. These contradictions provide the template for the general proof of theorem 1.1, exploiting the regularity conditions and the implied properties of the bid-shading behavior for any bidder.

First, suppose that $\epsilon_1 \equiv \inf\{x : F_X(x) \neq F^*(x)\} > 0$ and also suppose that $\epsilon_2 \equiv \inf\{x > \epsilon_1 : F_X(x) = F^*(x)\} > \epsilon_1$ so that for $y \in [0, \epsilon_1)$, $F_X(y) = F^*(y)$ and, as such, the inverse bidding functions are identical to one another in this region, i.e., $\sigma^{-1}(y) = \sigma_*^{-1}(y)$. Note that $\sigma^{-1}(y)$ is strictly greater than y away from the origin, continuous, and strictly increasing, so there is some $\tilde{y} < \epsilon_1$ with $\sigma^{-1}(\tilde{y}) \in (\epsilon_1, \epsilon_2)$. Then, $F_X(\sigma^{-1}(\tilde{y})) \neq F^*(\sigma_*^{-1}(\tilde{y}))$, contradicting 1.9. As such, since the bid-shade is non-negative, continuous, and zero at the origin, any candidate distribution over valuations satisfying the condition 1.9 must either be identical to the true distribution or differ from the true distribution starting at the origin.

Now, suppose the distributions F_X and F^* diverge immediately from the origin and, wlog, that $F_X(x) > F^*(x)$, for all $x \in (0, \epsilon)$, with the definition that $\epsilon \equiv \sup\{x : F_X(x) > F^*(x)\}$. In this case, the condition in equation 1.9 demands that $\sigma^{-1}(x) < \sigma_*^{-1}(x)$. However, there must come a point in $(0, \epsilon)$ where the distribution F^* begins "catching up" with F_X, i.e., where $f_X(x) < f^*(x)$. But since $F_X(x) > F^*(x)$, these inequalities imply $\frac{F_X(x)}{f_X(x)} > \frac{F^*(x)}{f^*(x)}$, contradicting the requirement that $\sigma^{-1}(x) < \sigma_*^{-1}(x)$. Here, monotonicity of bidding couples with the definition of F_X and F^* as the integral of f_X and f^*, respectively, to establish that any two distributions that diverge immediately from the origin cannot be observationally equivalent.

The following corollary summarizes the main identification result of this subsection:

Corollary 1.2 (Identification of $L1_T$ Bidder-Type) *Suppose that an econometrician observes the distribution over bids, $F_{S,L1_T}$, from a homogeneous population of $L1_T$ bidder-types, then the distribution over valuations, F_X, is identified.*

Random Level 2 ($L2_R$)

The identification argument for the Random Level 2 bidder-type closely mirrors the analysis of the Truthful Level 1 bidder-type since the Random Level 2 bidder-type best responds to the belief that other players' bids are a linear transformation of their valuation. Incorporating the known constant Jacobian term, Crawford and Iriberri (2007a) show the first order condition 1.8 above becomes:

$$X_i = \sigma_{L2_R}(X_i) + \frac{F_X\left(\frac{N}{N-1}\sigma_{L2_R}(X_i)\right)}{N f_X\left(\frac{N}{N-1}\sigma_{L2_R}(X_i)\right)}$$

The identification proof from the $L1_T$ bidder-type generalizes immediately, establishing identification:

Corollary 1.3 (Identification of $L2_R$ Bidder-Type) *Suppose that an econometrician observes sufficient data from a homogeneous population of $L2_R$ bidder-types to identify the distribution over bids, $F_{S,L2_R}$, then the distribution over valuations, F_X, is identified.*

Truthful Level 2 ($L2_T$)

Identification in the $L2_T$ case is complicated by the lack of a closed-form solution for the $L1_T$ bidding strategy. Crawford and Iriberri (2007a) characterize the first order condition 1.6 for the $L2_T$ bidder-type as:

$$\begin{aligned}(X-\sigma_{L2_T}(X))\, f_X\left(\sigma^{-1}_{L1_T}(\sigma_{L2_T}(X))\right)\frac{d\sigma^{-1}_{L1_T}(s)}{ds}|_{s=\sigma_{L2_T}(X)}\\ -F_X\left(\sigma^{-1}_{L1_T}(\sigma_{L2_T}(X))\right)=0\end{aligned} \tag{1.10}$$

While no closed-form solution exists for the $L1_T$ bidding function, equation 1.8 gives the inverse of the $L1_T$ bidding function, with corresponding derivative:

$$\frac{d\sigma^{-1}_{L1_T}(s)}{ds}=1+\frac{f_x(s)^2-F_X(s)\, f'_X(s)}{(N-1)\, f_X(s)^2}=\frac{N}{N-1}-\frac{F_X(s)\, f'_X(s)}{(N-1)\, f_X(s)^2}$$

Substitute this identity into equation 1.10 and rearranging gives the $L2_T$ inverse bidding function:

$$\begin{aligned}X_i=&\sigma_{L2_T}(X_i)\\ &+\frac{F_X\left(\sigma_{L2_T}(X_i)+\frac{F_X\left(\sigma_{L2_T}(X_i)\right)}{(N-1)f_x\left(\sigma_{L2_T}(X_i)\right)}\right)}{f_X\left(\sigma_{L2_T}(X_i)+\frac{F_X\left(\sigma_{L2_T}(X_i)\right)}{(N-1)f_x\left(\sigma_{L2_T}(X_i)\right)}\right)\left(N-\frac{F_X\left(\sigma_{L2_T}(X_i)\right)f'_X\left(\sigma_{L2_T}(X_i)\right)}{f_X\left(\sigma_{L2_T}(X_i)\right)^2}\right)}\end{aligned} \tag{1.11}$$

Here, the distribution over bids for the $L2_T$ bidder-type also depends on the derivative of the pdf for true valuations, introducing a new potential source for confounding identification. However, given condition 4 in Theorem 1.1, the bracketed expression in the denominator of equation 1.11 is positive and bounded away from zero, with implicit differentiation establishing the bidding equation $\sigma_{L2_T}(X_i)$ as monotonic in valuation X_i.

The identified set is characterized by the consistency requirement in equation 1.9 from section 1.4.2 applied to the first order conditions in 1.11. Here the argument for identification has little structural difference, except the requisite continuity conditions apply to the higher order derivatives of the $L1_T$ bidder-type's strategy. Nonetheless, the approach of establishing contradictions through analyzing the bid shade is effectively unchanged.

Corollary 1.4 (Identification of $L2_T$ Bidder-Type) *Suppose that an* *econometrician observes distribution over bids, $F_{S,L2_T}$, from a homogeneous population of $L2_T$ bidder-types, then the distribution over valuations, F_X, is identified.*

General Level-k

The general level-$(k-1)$ bidding strategy is a continuous function of the bidder's signal and $(k-1)$ derivatives of the pdf over valuations, hence each iteration of the cognitive hierarchy requires another continuous derivative of the distribution over valuations as indicated in assumption L1.1.3 in lemma 1.1. The regularity conditions in lemma 1.1 ensure this bidding strategy has derivatives that exist, are bounded, and continuous, giving rise to a continuous, monotonic level-k bidding function. The general level-k first order conditions from Crawford and Iriberri (2007a) are:

$$\begin{aligned}\left(X-\sigma_{Lk_\tau}(X)\right)(N-1)\,f_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right)\frac{d\sigma^{-1}_{L(k-1)_\tau}(s)}{ds}\Big|_{s=\sigma_{Lk_\tau}(X)}\\ =F_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right)\end{aligned}$$

Rearranging this equation gives the inverse bidding function that characterizes consistency required for the distribution over bids to be generated by the distribution over valuations:

$$X=\sigma_{Lk_\tau}(X)+\frac{F_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right)}{(N-1)\,f_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right)\frac{d\sigma^{-1}_{L(k-1)_\tau}(s)}{ds}\Big|_{s=\sigma_{Lk_\tau}(X)}} \qquad (1.12)$$

The observational equivalence arguments developed in previous sections are applied to the relationship in equation 1.12 in the proof in the appendix.

1.4.3 Testability of the Level-k Behavioral Model

Given identification, the next step is to characterize overidentifying restrictions that could reject the level-k behavioral model as mis-specified. As is the case with the equilibrium behavioral model, if the econometrician always observes auctions with a fixed number of participating bidders (for example, if every auction has exactly 5 bidders), the level-k behavioral model imposes no testable restrictions on the data beyond those implied by independence of individual bidding decisions and those required for identification. That is, for any distribution over bids, $\hat{F}_S$, there exists a corresponding distribution over valuations, $\hat{F}_{X,Lk_\tau}$, admitting a strictly monotonic inverse bidding function that is consistent with the hypothesis that the entire population of bidders is of type Lk_τ. This result follows from the identification results above, each of which hold independently for any distribution over bids.

Again in parallel to equilibrium results, the level-k model is testable if the number of participating bidders in the auction varies exogenously, (for example, if half of the auctions in the sample have 5 competing bidders and half have 20 bidders). The level-k behavioral model defines precisely how a bidder's strategy reacts to changes in the number of bidders participating in an auction, imposing a continuum of over-identifying restrictions, one for each quantile of the distribution over valuations. Suppose the econometrician observes just two distributions over bids corresponding to two different levels of competition in the auction, $\hat{F}_{S_{N_1}}$ and $\hat{F}_{S_{N_2}}$ and wishes to test the hypothesis that the entire population of bidders is of type Lk_τ. The econometrician can then use the two distributions over bids to recover two distribution over valuations, $\hat{F}_{X,Lk_\tau,N_1}$ and $\hat{F}_{X,Lk_\tau,N_2}$. If the hypothesized behavioral model is true, these two recovered distributions must be equal, that is, $\hat{F}_{X,Lk_\tau,N_1}(x) = \hat{F}_{X,Lk_\tau,N_2}(x)$, for all x. If any of the quantiles from the two distributions disagree, then the hypothesized behavioral model can be rejected. This testability result applies not only to hypotheses that the population is homogeneous, but can also be used to test hypotheses about mixtures of the populations, providing the basis for the identification argument in heterogeneous populations.

In another parallel to results from Athey and Haile (2002) and Haile, Hong, and Shum (2003), given observation of more than one bid from each auction, these different order statistics from the distribution over bids can test implications of the IPV model. For example, given the winning bid, denoted S_w and the second highest bid S_l, the econometrician could estimate the distribution over valuations from both samples of data individually and test the hypothesis that the estimated distributions are equivalent up to sampling error. However, it is important to note that this testable implication derives primarily from the independence of private values and, in particular, of information and beliefs, which implies independence of observed bids. In this case, the distributions over bids recovered from S_w and S_l as the first and second order statistics for i.i.d. random variables should be identical. As such, while the restriction would be violated in the affiliated values problem or in the presence of unobserved heterogeneity, it has no power against either behavioral mis-specification.

1.5 Identification in Heterogeneous Populations

In heterogeneous populations, the need to identify the distribution over bidder-types in addition to the distribution over valuations introduces a free dimension to the model that requires additional information for identification. One such information set allowing identification corresponds to cases in which the econometrician can estimate each individual's distribution over bids. While this setting is unlikely to obtain in application, it illustrates the logic for separating bidder-types through the relationships among their distributions over bids. When the econometrician does not repeatedly observe individual bidders, the model is incompletely identified as the econometrician can construct a distribution over valuations that generates the observed distribution over bids for any fixed distribution over bidder-types. In this case, identification relies on the testable implications in the previous section to establish exclusion restrictions that recover identification when the benchmark informational setting is augmented with exogenous variation in the number of bidders.

The expanded model with heterogeneous bidder types requires some notation regarding the population distribution over bidder-types and an independence assumption for bidder-type assignment. Define the set of K possible bidder-types by $\mathcal{K}$ and denote the distribution over bidder-types, $p = [p_1, \ldots, p_K]$, so that the probability that a bidder is of type k is p_k. Further, assume the assignment of bidder-types is independent of that individual's valuation.

Assumption 1.3 (Independent Assignment to Bidder-Types)
Each player $i \in N$ is randomly assigned a unique bidder-type $\tau(i) \in \mathcal{K}$ according to the distribution $p = [p_1, \ldots, p_K]$ independently of the number of bidders in the auction and the player's latent valuation.

The econometric structure now consists of the true distribution over valuations, F_X, the set of behavioral types, $\mathcal{K}$, and the distribution over behavioral types, p. Since the set of behavioral types is defined by the economic theory, treat $\mathcal{K}$ as known. The generalized definition of observational equivalence in heterogeneous populations becomes:

Definition 1.2 (Level-k Observational Equivalence)
Given the set of bidder types, a structure $(F_X, \mathcal{K}, p)$ is observationally equivalent to the structure $(F_*, \mathcal{K}, p_*)$ if:

$$\sum_{k=1}^{K} p_{X,k}(s) F_X\left(\sigma_{k,X}^{-1}(s)\right) = F_S(s) = \sum_{k=1}^{K} p_{*,k}(s) F_*\left(\sigma_{k,*}^{-1}(s)\right)$$

where, $p_{\cdot,k}(s) = \frac{p_{\cdot,k} F_\cdot\left(\sigma_{k,\cdot}^{-1}(s)\right)}{\sum_{\kappa=1}^{K} p_{\cdot,\kappa} F_\cdot\left(\sigma_{\kappa,\cdot}^{-1}(s)\right)}$

1.5.1 Identification from Repeated Individual Observation

Suppose the population of bidder-types is constant in each auction and the econometrician observes each individual's bidding behavior across a large number of independent auctions. This information is sufficient to characterize each individual's distribution over bids and, consequently, separate all bidders into K groups of bidder-types. All that remains to establish identification in this setting is to uniquely sort each of the observed bidder-types to a position in the behavioral hierarchy.

In the case $\mathcal{K} = \{L0_T, L1_T, L2_T, L0_R, L1_R, L2_R, Eqm\}$, analyzed in Crawford and Iriberri (2007a), the $L0_T$ bidder-type as the bidder with the distribution over bids that first-order stochastically dominates all other bidder types' distributions. This information is sufficient to identify the distribution over valuations, which in turn identify the bidder-types corresponding to each of the other distributions over bids by computing the implied distribution over bids for each bidder-type and matching it to the population distribution over bids. If there are no truthful bidders, the $L1_R$ bidder type's linear strategy identifies this bidder-type's distribution as the bidder with the thinnest right tail. This result is due to Battigalli and Siniscalchi (2003), who show that best responding to iteratively rationalizable bidding functions yields bidding functions that are generally concave and weakly decreasing in the iterations.[6] Once a single distribution over bids is assigned to the bidder-type generating that distribution, the distribution over valuations is identified based on this association, an algorithm that applies due to the finite number of bidder-types that is not available to Aradillas-Lopez and Tamer (2008). This result is summarized in the next theorem:

Theorem 1.2 (Identification with Repeated Individual Observation)
Suppose:

1. *The econometrician observes distribution over bids for each individual in a heterogeneous population of level-k bidders responding to a total population of N possible bidders, and*

2. *The distribution over valuations is identified for each bidder-type $k \in \mathcal{K}$ from the distribution over bids, $F_{S,k}$, observed in homogeneous populations,*

then the distribution F_X is uniquely identified.
Suppose further that:

3. *The econometrician observes bidding behavior for a homogeneous population of level-k bidders responding to a total population of $N^* \neq N$ possible bidders in each auction*

then the level-k behavioral model and the distribution over bidder-types are jointly testable.

[6]In models with a unique Level-0 bidder-type, the set of recovered distributions over bids can be sorted into bidder-types based on their upper support. In particular, if the level $k-1$ bidder-type never bids above $\overline{s}_{k-1}$, then the level k bidder's support will be weakly smaller than $\overline{s}_{k-1}$.

Proof. The information set for the econometrician consists of the K distributions over bids, $F_{S,\tau_1}, \ldots, F_{S,\tau_K}$. The identification task is then to assign these types $(\tau_1, \ldots, \tau_K)$ to one of the $K!$ possible permutations of true types. The restrictions that accomplish this task are generated by the fact that each of the behavioral types draw their valuations from the same distribution over valuations.

Suppose τ_1 is known, then the distribution over valuations $F_{X_{\tau_1}}$ is identified from F_{S,τ_1}. For the true value of τ_2 and any bid value s:

$$F_{S,\tau_2}(s) = F_{X_{\tau_1}}\left(\sigma_{\tau_2}^{-1}(s)\right) = F_{S,\tau_1}\left(\sigma_{\tau_1}\left(\sigma_{\tau_2}^{-1}(s)\right)\right) \tag{1.13}$$

Identification would fail if there were two values of τ_2 that could satisfy equation 1.13. However, if this condition were satisfied by two different types, it would imply their distribution over bids were the same so that their bidding functions are identical, in which case one of the two types is redundant in the behavioral specification. It may also be the case that there does not exist a compatible sort, which would reject the level-k model as mis-specified.

The testable implications for the cognitive hierarchy model with variation in the number of bidders are analogous to those established in Athey and Haile (2002) and discussed in section 1.4.3. As the number of bidders changes, the bidder-type's strategies change deterministically as a function of the distribution over bids and the number of bidders. As such, having estimated F_X in a setting with N_1 bidders with distribution over bids the distribution over bids for N_2 bidders, F_{S,N_2}, is completely determined as long as F_X itself doesn't depend on the number of bidders (which commonly obtains in practice). Every quantile of F_{S,N_2} provides a testable restriction of the level-k model. ■

1.5.2 Identification under Pooled Bidding Behavior

The econometrician rarely observes the information set studied in the previous subsection. In the rare cases that individual bidding data is obtainable (for example, in sealed-bid auctions), anonymity concerns typically prevent tracking an individual bidder across auctions and the repeated interactions among bidders is likely to warrant additional strategic analysis. More commonly, the econometrician is capable of observing the bids of all individuals in the population without being able to follow them from one auction to the other. In this benchmark informational setting, where the econometrician only observes sufficient information to identify the population distribution over bids, the model is incompletely identified. For example, the econometrician generally does not have sufficient information to differentiate whether that distribution over bids was generated by a population consisting entirely of truthful bidder-types or entirely of $L1_R$ bidder-types. The next theorem establishes the incomplete identification result as even more severe in that, for any given distribution over bids and any given distribution over types, there exists a distribution over valuations that generates the observed distribution over bids.

Theorem 1.3 (Partial Identification in Heterogeneous Populations) *Suppose the econometrician observes the distribution over bids in a fixed population of N bidders is $F_{S_N}(x)$ and, further, that the distribution over valuations is identified for each bidder-type $k \in \mathcal{K}$ from the distribution over bids, F_{S_N} in homogeneous populations. In this case, for any distribution over behavioral types, $[p_1, \ldots, p_K]$, there exists a distribution $F_X(x)$ generating $F_{S_N}(x)$.*

The proof in the appendix exploits the structure imposed by equation 1.13 to separate the mixture distribution $F_{S_N}(x)$ into its component distributions over bids for homogeneous populations. Having recovered the implied distributions of bids from homogeneous bidder-types, Theorem 1.2 proves that there exists a unique distribution over valuations that generates the recovered distributions over bids. The requirement that the distribution over valuations is constant across bidder-types is enforced through the decomposition of the mixture distribution over bids in the first step. The unique distribution over latent valuations for any and every hypothesized mixture of bidder-types establishes the partial identification result.

Figure 1.4 illustrates this result in the simple setting with $\mathcal{K} = \{L0_T, L1_R\}$, $p_{L0_T} = 0.7$, and $p_{L1_R} = 0.3$ when the true valuations are exponentially distributed. Without knowing p, the true distribution of valuations could correspond to any one of the cdf's in the figure, with the z-axis providing depth to indicate the mixture of bidder-types that generates the distribution over bids from the hypothesized distribution over valuations. While using the correct distribution of bidder-types recovers the true distribution of valuations, other distributions of bidder-types are estimated under mis-specified behavioral models, leading to incorrect inference. Since the true distribution over behavioral bidder-types is unidentified, the behavioral model only characterizes the identified set of observationally equivalent distributions over valuations, which maps into the unit symplex defined by the set of distributions over types.

To recover identification, suppose the econometrician observes bidding distributions when either 5 or 20 bidders participate in the auction, where all bidder-types are drawn from a constant distribution over bidder-types, p. In this case, only the true distribution over bidder-types will identify the same distribution over valuations almost everywhere from both of the distributions over bids observed in the two different population sizes. further, in settings where the level-k behavioral model does not hold, the continuum of testable restrictions due to variation in the number of bidders provides enough information to reject the model if the intersection of the identified sets for the distribution over valuations in the two populations is empty.

Figure 1.5 presents the examples from Figure 1.4 when the distribution over bids is created under the same true distribution of valuations but with $N = 20$ bidders. As illustrated in Figure 1.6, the intersection of the identified set for $N = 5$ with the identified set for $N = 20$ occurs at the true distribution over bidder types, selecting a unique estimated distribution over valuations that remains stable with changes in the number of bidders. If, however, variation in the number of bidders also results in variation in the distribution over bidder-types or shifts in the distribution over valuations, then the partial identification result from Theorem 1.3 applies. The following theorem summarizes this result:

Theorem 1.4 (Identification with Variation in Number of Bidders)
Suppose:

1. *The econometrician observes the distributions over bids, $F_{S_{N_1}}$ and $F_{S_{N_2}}$, for a heterogeneous population of level-k bidders responding to a total population* *of N_1 and N_2 possible bidders, with $N_1 \neq N_2$,*

2. *The distribution over valuations is identified for each bidder-type $k \in \mathcal{K}$ from the distribution over bids in homogeneous populations, and*

3. *For any $k_1, k_2 \in \mathcal{K}$ with $k_1 \neq k_2$, the set:*
   $\{x \in [\underline{x}, \overline{x}] \, | \sigma_{k_1,N_1}(x) - \sigma_{k_1,N_2}(x) \neq \sigma_{k_2,N_1}(x) - \sigma_{k_2,N_2}\}$
   has nonzero Lebesgue measure whenever $i \neq j$.

then both the distribution over valuations, F_X, and the distribution over behavioral types, p, are uniquely identified. Further, the level-k behavioral model is testable through overidentifying restrictions.

1.6 Maximum Likelihood Estimation

To begin, assume the distribution over valuations is drawn from a compact family of distributions indexed by the infinite-dimensional parameter vector θ. Characterizing the expected log-likelihood criteria function begins by analyzing the data generating process, stating the likelihood of the data in cases where the econometrician observes all bids or possibly only observes the winning bid in an auction. Here the heterogeneous behavioral types within the structural model gives these likelihoods a non-standard mixture structure. Consistent estimation follows upon adopting the sieve space proposed by Bierens (2006) and Bierens and Song (2007) using a strong uniform law of large numbers due to Artstein and Wets (1995) to address discontinuities in the upper semicontinuous criterion function.

1.6.1 Parametric Likelihood Functions

Suppose bidder i is of the kth bidder-type, so the observed bid, $s_i = \sigma_k(x_i)$ and, equivalently, $x_i = \sigma_k^{-1}(s_i)$. Then the cumulative likelihood of having observed a bid less than s_i, conditional on the true parameter vector θ is:

$$F_{S,k}(s_i;\theta) = F_X\left(\sigma_k^{-1}(s_i;\theta);\theta\right)$$

Differentiating with respect to s_i and applying the Jacobian of the inverse bidding function, the likelihood of observing a bid equal to s_i for bidder-type k is:

$$f_{S,k}\left(s_i;\theta\right)=\frac{f_X\left(\sigma_k^{-1}\left(s_i;\theta\right);\theta\right)}{\sigma_k'\left(\sigma_k^{-1}\left(s_i;\theta\right);\theta\right)}$$

To characterize this mixture structure, define an indicator that the ith bid was chosen by a k-type bidder by $d_{ik}=1_{\{\tau(i)=k\}}$, which can be stacked into the vector $d_i=\left(d_{i1},\ldots,d_{iK}\right)'$. The likelihood for the ith observation conditional on the ith bidder's type can then be stated in either of two forms:

$$f\left(s_i;d_i,\theta\right)=\sum_{k=1}^{K}d_{ik}f_{S,k}\left(s_i;\theta\right)=\prod_{k=1}^{K}f_{S,k}\left(s_i;\theta\right)^{d_{ik}}$$

Since the bidder's type is independent of their valuation, the distribution for the type generating the ith bid is a multinomial random variable with distribution:

$$p_k\left(\theta\right)=Pr\left(\tau\left(i\right)=k\right)=\prod_{k=1}^{K}p_k^{d_{ik}}$$

Combining these two results gives the distribution of the ith bid conditional on the distribution over valuations:

$$f\left(s_i;\theta\right)=\prod_{k=1}^{K}p_k^{d_{ik}}f_{S,k}\left(s_i;\theta\right)^{d_{ik}} \tag{1.14}$$

The unconditional likelihood of observing all bids then provides the basis for the expected log likelihood that serves as the criterion function for maximum likelihood estimation:

$$\mathcal{L}_T\left(\theta;s_1,\ldots,s_T\right)=\prod_{i=1}^{T}\prod_{k=1}^{K}p_k^{d_{ik}}f_{S,k}\left(s_i;\theta\right)^{d_{ik}} \tag{1.15}$$

$$\hat{\Psi}_T\left(\theta;s_1,\ldots,s_T\right)=\sum_{i=1}^{T}\sum_{k=1}^{K}d_{ik}\ln p_k+d_{ik}\ln f_{S,k}\left(s_i;\theta\right) \tag{1.16}$$

In a panel sample with repeated observations of an individual following a constant bidding strategy, the likelihood has additional structure reflecting the additional information about that bidder's type. Denoting the sample of T_i bids for individual i by $S_i = \{S_{i,1}, \ldots, S_{i,T_i}\}$, the probability of observing a sample of bids $s_i = \{s_{i,1}, \ldots, s_{i,T_i}\}$ given $\tau(i) = k$ is:

$$f_{S,k}(s_i;\theta) = \prod_{t=1}^{T_i} \frac{f_X\left(\sigma_k^{-1}(s_i;\theta);\theta\right)}{\sigma_k'\left(\sigma_k^{-1}(s_i;\theta);\theta\right)}$$

Since this is the only basic definition that changes, the likelihood for the full bidding sample remains as stated in equation 1.15.

In the benchmark setting, the econometrician only observes the winning bid, as in a Dutch descending auction where the auction ends once the winning bidder claims the object at the announced price. The distribution of winning bids is given by the distribution for the maximum bid, which will depend on the actual mixture of types in each round. As such, computing the unconditional distribution for winning bids requires summing over all possible mixtures of the K bidder-types. Defining $F_{S_N,k}$ as the distribution over bids for the k-th bidder-type in an auction with N participating bidders, the resulting distribution is most readily stated in terms of cumulative densities:

$$F_{W_N}(w;\theta) = \qquad (1.17)$$

$$\sum_{n_1=0}^{N} \sum_{n_2=0}^{N-n_1} \cdots \sum_{n_{K-1}=0}^{N-\sum_{k=1}^{K-2} n_k} \binom{N}{n_1}\binom{N-n_1}{n_2} \cdots \binom{N-\sum_{k=1}^{K-2} n_k}{n_{K-1}} \prod_{k=1}^{K} p_k^{n_k} F_{S_N,k}(w;\theta)^{n_k}$$

Note that the distribution over winning bids is a continuous polynomial in the distribution over bidder-types implying the expected value of the winning bid, or the expected revenue from the auction, is also continuous in the distribution over bidder-types. The distribution in 1.17 plays a central role in characterizing how the optimal reserve price is affected by the distribution over bidder-types.

1.6.2 Semi-Nonparametric Consistent Estimation

The parametric distributional analysis defines a semi-nonparametric maximum likelihood (SNP-ML) sieve-based estimator, with consistency following from techniques surveyed in Chen (2007). Consistency requires four key conditions: Identification, Compact Parameter Space, A Uniform Strong Law of Large Numbers (USLLN) for the finite-sample Criterion function, and a truncation algorithm. Sections 1.4 and 1.5 establish identification, so there is a unique maximum (θ_0, p_0) to the population criterion function:

$$\Psi(p, \theta) = E\left[d_k \ln p_k + d_k \ln f_{S,k}(S;\theta) \middle| p, \theta\right] \tag{1.18}$$

This subsection begins with a review the Legendre polynomial sieve space and compactness results from Bierens (2006) before introducing the USLLN for the criterion function in equation 1.18 and verifying consistency for SNP-ML estimator. The consistency argument closely follows Bierens (2006)'s analysis of interval-censored mixed proportional hazard models, but is complicated by the upper semicontinuous log-likelihood objective function. As such, neither the USLLN in Bierens (2006)'s initial treatment nor that in Bierens and Song (2007)'s generalization of Jennrich's USLLN apply. However, adapting the USLLN from Artstein and Wets (1995), who use the notion of weak epi-convergence to establish uniform convergence of functions, is a straightforward exercise.

Legendre Polynomial Sieve Space

Consistency for extremum estimators typically requires the parameter space to be a compact metric space to avoid measurability problems. Further, sieve spaces require a set of orthogonal basis functions that avoid the ill-posedness problem in which a single function may have multiple (almost sure) equivalent representations within a truncated set of basis functions. Bierens (2006) and Bierens and Song (2007) present such a space based on the Legendre polynomials, which represent an orthonormal basis for the set of functions on the unit interval. To address the range of the sieve space, Bierens and Song (2007) map the support of the distribution over valuations to the unit interval using an absolutely continuous distribution function such as the exponential cumulative density function with unbounded support. Embedding the support for valuations in the unit interval addresses measurability for the polynomial sieve by bounding the arguments of the density function. Given the invertible mapping $G : [0, \infty] \to [0, 1]$, there is some distribution H over the unit interval so that the cumulative distribution and probability density functions for valuations as:

$$\begin{aligned} F_X(x) &= H(G(x)) \\ f_X(x) &= h(G(x))\, g(x) \end{aligned} \tag{1.19}$$

For computational purposes, instead of using a non-linear transformation of the valuations, the following assumption embeds the set of possible valuations in the unit interval using a linear transformation:

Assumption 1.4 (A Priori Bounded Support) *There exists upper and lower bounds for valuations $\overline{M} > \overline{x} > \underline{x} > \underline{M}$ known to the econometrician a priori. The support for valuations is mapped onto the unit interval using the linear transformation:*

$$\begin{aligned} G(x) &= \frac{x - \underline{M}}{\overline{M} - \underline{M}} \\ g(x) &= \frac{1}{\overline{M} - \underline{M}} \end{aligned}$$

In practice, these bounds have little effect on the estimation process, though $\underline{M}$ should be less than the minimum observed bid. Having mapped the support for valuations into the unit interval, Bierens' and Bierens & Song's sieve space based on constrained Legendre polynomials can approximate the distribution H. The linear mapping into the unit interval coupled with the Legendre polynomial sieve facilitates computation of the maximum likelihood estimator by admitting analytical solutions for a number of the formula that enter into calculations of the equilibrium bidding function. The detailed representation of the sieve space and results establishing compactness appear in Appendix 1.A.2 for interested readers.

Maximum Likelihood Estimator Consistency

Consistency requires verifying three key convergence properties for finite-sample criterion functions. The uniform strong law of large numbers for the criterion function is a special case of the general uniform strong law of large numbers in Artstein and Wets (1995). Convergence of the criterion function at its optimum follows from the USLLN and continuity of the criterion function at the optimum. Lastly, Bierens and Song (2007) show that the criteria function's optimum over a sequence of constrained sieve spaces converges to the the global optimum of the unconstrained sieve space. These results are summarized in the following lemma.

Lemma 1.3 (Upper Semicontinuous Random Function Convergence)

a. *Let Θ be a compact metric space with metric $\rho(\theta_1, \theta_2)$, and let $\Psi_t(\theta), t = 1, 2, \ldots, T, \ldots$ be a sequence of i.i.d. random, real valued, upper semicontinuous functions on Θ. If, in addition, for each $\theta_0 \in \Theta$, there exists an open set $Q_0 \subset \Theta$ and a constant $\xi_0 < \infty$ such that*

$$\sup_{\theta \in Q_0} \Psi_1(\theta) < \xi_0 \ \ a.s.$$

then

$$\sup_{\theta \in \Theta} \left| \frac{1}{N} \sum_{j=1}^{N} \Psi_j(\theta) - \Psi(\theta) \right| \to 0 \ \ a.s.$$

b. *Suppose further that $\Psi(\theta)$ is an upper semicontinuous real function on Θ, define $\hat{\Psi}_N(\theta) = \frac{1}{N} \sum_{j=1}^{N} \Psi_j(\theta)$, and let $\hat{\theta}_N = \arg\max_{\theta \in \Theta} \hat{\Psi}_N(\theta)$ and $\theta_0 = \arg\max_{\theta \in \Theta} \Psi(\theta)$. Then for $N \to \infty$,*

$$\Psi\left(\hat{\theta}_N\right) \to \Psi(\theta_0) \ \ a.s.$$

If θ_0 is unique, then $\rho\left(\hat{\theta}_N, \theta_0\right) \to 0$ a.s.

c. *Let $\{\Theta_n\}_{n=0}^{\infty}$ be an increasing sequence of compact subspaces of Θ for which the computation of*

$$\hat{\theta}_{n,N} = \arg\max_{\theta \in \Theta_n} \hat{\Psi}_N(\theta) \tag{1.20}$$

is feasible. Suppose that for each $\theta \in \Theta$ *there exists a sequence* $\theta_n \in \Theta_n$ *such that* $\lim_{N\to\infty} n_N = \infty$*, and denote the sieve estimator involved by* $\tilde{\theta}_N = \hat{\theta}_{n,N}$*. Then* $\rho\left(\tilde{\theta}_N, \theta_0\right) \to 0$ *a.s.*

Proof. The conditions in statement (a) are strictly stronger than the sufficient conditions for the uniform strong law of large numbers in Artstein and Wets (1995) Theorem 2.3 but are easily verified to apply to the level-k model. The remaining results follow immediately from combining this strong law of large numbers with the arguments in Bierens and Song (2007), Theorems (1) - (3). ■

These convergence results provide the basis for SNP-consistent estimation of the level-k auction model. Given compactness results for the Legendre polynomial sieves in lemma 1.4, let the distribution over valuations be indexed by H, the equivalent distribution over the unit interval from equation 1.19 that admits a Legendre polynomial representation. Redefine $\theta = [p', H]'$ to join the distribution over types and distribution over valuations into a single parameter vector belonging to a metric space, Θ, with the metric:

$$\rho\left(\theta_1, \theta_2\right) = \max\left[\max\left|p_1 - p_2\right|, \sup_{0\leq u\leq 1}\left|H_1\left(u\right) - H_2\left(u\right)\right|\right] \tag{1.21}$$

The population criterion function is given by $\Psi\left(\theta\right)$ from equation 1.18, with the sample counterpart $\hat{\Psi}_T\left(\theta\right)$ from equation 1.16.

Since the valuations are bounded above and all players follow continuously differentiable, strictly monotonic strategies, there are no atoms in the distribution over bids and the criterion function satisfies the local uniform bound required for the uniform strong law of large numbers in lemma 1.3, part a. The upper semi-continuity in lemma 1.3, part b, follows by similar logic, with discontinuities due to bidder-types having different supports for their distributions over bids. Finally, the convergence properties required for lemma 1.3, part c, are satisfied given the conditions for lemma 1.4 in Appendix 1.A.2 where $\Theta_N = \{\theta \in \Theta | \delta_{N+j} = 0, j = 1, 2, \ldots\}$. The identification results of the previous sections establish the uniqueness of the optimum to the population criterion function, giving the following consistency result:

Theorem 1.5 (Consistency of SNP-ML Estimator)

Suppose Assumptions 1.1 - 1.4 hold, Theorems 1.1, 1.2, or 1.4 apply so that the level-k auction model is identifed, and Lemma 1.4 applies. Let n_N *be an arbitrary subsequence of* n *such that* $\lim_{N\to\infty} n_N = \infty$*, then for the estimator* $\tilde{\theta}_N$ *defined in equation 1.20,* $\rho\left(\tilde{\theta}_N, \theta_0\right) \to 0$ *a.s.*

1.6.3 Introducing Auction-Specific Covariates

The experimental ideal informational setting for estimating the auction model would be to observe bidding behavior in a series of identical auctions. While these settings do not obtain in empirical work, a common practice controls for auction-specific heterogeneity by allowing the distribution over valuations to depend upon a set of observable auction-specific covariates. In fact, the link between observable features of the object at sale and the distribution over valuations is frequently the primary concern of the empirical exercise.[7] For additional tractability, assume a separable structure for individual valuations so that:[8]

Assumption 1.5 (Separable Auction-Specific Heterogeneity) *Given a set of auction-specific covariates,* Z_t*, bidder i's valuation for the object at auction,* X_{it} *is given by:*

$$log\left((X_{it}\right) = \gamma' Z_t + U_{it}$$

where $U_{it}|Z_t \perp\!\!\!\perp U_{jt}|Z_t, \forall i \neq j$.

The linear separable form of auction-specific heterogeneity can be relaxed and is made here for numerical and notational parsimony that allows simply appending γ to the parameter vector θ. However, the introduction of covariates does

[7] Note that I do not address unobserved heterogeneity across auctions here. A great deal of research had focused on developing strategies for addressing auction-specific heterogeneity in estimating auctions. A number of works, including Bajari and Ye (2003) and Hong and Shum (2002) link unobserved heterogeneity tot he number of bidders in the auction. Haile, Hong, and Shum (2003) use multiple bids observed in each auction to control for auction specific heterogeneity, a strategy further developed by Krasnokutskaya (Forthcoming - 2009), An, Hu, and Shum (2009), and Hu, McAdams, and Shum (2009). It is likely that a similar technique could be adapted to level-k auctions but such an analysis is beyond the scope the current work.

[8] While this is a relatively restrictive assumption, some novel tests have recently been developed by Lu (2009) using generic characteristic revealing functions associated with tests for conditional independence that allow for testing this restriction.

slightly affect consistenty since the support for the distribution over bids now depends on the parameters themselves, introducing a discontinuity at the optimum of the finite-sample objective function. The approach from Donald and Paarsch (1996) addresses this issue by modeling the relationship between the support for bids and parameters of interest through a set of additional constraints that adapt naturally to the sieve specification. For exposition, assume the lower support of the distribution over bids is zero and define the upper support for the distribution over bids conditional on covariates $Z_t = z \in \mathcal{Z}$ and parameter vector θ:

$$\overline{s}\left(\theta, z\right) = \max_{k \in \mathcal{K}} \sigma_k \left(\overline{x}; \theta, z\right)$$

To ensure all observed bids fall in the support of the distribution over bids conditional on auction covariates and the parameter vector, restrict the finite-sample parameter space to:

$$\Theta_{N,T} = \{\theta \in \Theta_N | 0 \leq S_t \leq \overline{s}\left(\theta_0, z\right) t = 1, 2, \ldots, N\} \tag{1.22}$$

Lastly, additional assumptions ensure that the convergence result from Lemma 1.3, part c, apply and the retrictions in equation 1.22 are not binding asymptotically. As in Donald & Paarsch's analysis, these assumptions take the form of a continuity condition on the upper support of bids and a restriction of the behavior for the distribution over bids near its upper support. From Paarsch and Hong (2006):

Assumption 1.6 (Restriction on Bid Distribution Upper Support)

a. *For any $\theta \in \Theta$, $\overline{s}\left(\theta, z\right)$ is continuous in z on $\mathcal{Z}$, and,*

$$\underline{x} < \inf_{z \in \mathcal{Z}} \overline{s}\left(\theta, z\right) < \sup_{z \in \mathcal{Z}} \overline{s}\left(\theta, z\right) < \infty$$

b. *For any $\epsilon > 0$,*

$$\inf_{z \in \mathcal{Z}} Pr\left[S > \overline{s}\left(\theta_0, z\right) - \epsilon\right] = \eta\left(\epsilon\right) > 0$$

Combining Paarsch and Hong (2006)'s Theorem 4.3.1 with Bierens and Song (2007)'s Theorem 3 establishes the convergence of the truncated sieve space to the limiting sieve space, with the proof of theorem 1.5 following immediately.

1.6.4 Addressing Computational Challenges

The level-k auction model model results in two key computational challenges. First, simply computing the likelihood requires solving for extremely non-linear bidding functions and inverse bidding functions that preclude analytical solutions. Second, the mixture-of-types structure of the problem requires an expectation-maximization (EM) algorithm to maximize the likelihood. The standard EM algorithm is computationally infeasible given the challenges associated with maximizing the expected likelihood and so I propose a feasible generalized EM algorithm with much faster convergence properties.

The primary computational challenge in implementing maximum likelihood methods for the equilibrium auction model is computing the equilibrium bidding and inverse bidding function. While the earlier equilibrium identification analysis worked directly with first-order conditions, it is well known that the equilibrium bidding function takes the form:

$$\sigma(x) = x - \frac{\int_{\underline{x}}^{x} F_X(u)^{N-1}\, du}{F_X(x)^{N-1}} \tag{1.23}$$

The integral here presents the crux of the computational challenge, particularly when inverting the equilibrium bidding function. For instance, if F_X is a 5-th order polynomial and $N = 11$, the analytical solution to the integral would require computing an 5^{10}-th order polynomial. As such, even though the analytical solution exists, numerical stability precludes its calculation in the presence of a large number of bidders or when the order of the polynomial sieve becomes large. Further, as the order of the polynomial grows, the number of terms in the analytical solution grows to the point where calculation becomes infeasible. Numerical stability requires limiting the data sample to focus on auctions with a relatively small number of bidders, as there is simply no other way to limit the maximal order of the polynomial. Once the order of the polynomial becomes sufficiently large, then computing the integral using quadrature provides a faster computation while retaining remarkable precision. Quadrature is especially appealing when there is a large number of bids that need to be inverted, as this allows partitioning the integral into very small segments.

Due to the latent mixture model in level-k auctions, the log-likelihood objective function is not directly observable. In standard mixture models, the expectation maximization (EM) algorithm addresses this nonobservability by treating the latent type-class as an unobserved variable, maximizing the expected likelihood through a series of improved approximations to the unobserved likelihood. This algorithm is computationally demanding even in settings where the arguments that maximize the likelihood can be calculated through closed form solutions. In the present application, analytical solutions exist for the distribution over types that optimizes the expected likelihood conditional on the distribution over valuations, but no analytical solutions exist for for the likelihood-maximizing distributional parameters, requiring a cumbersome numerical optimization procedure.

To address this issue, I propose a generalized expectation maximization (GEM) algorithm that partitions the likelihood maximization problem into two sub-problems: one for which quickly computable analytical solutions for the optima are available and one that requires a very slow numerical optimization step. This novel GEM algorithm maintains convergence to the optimum of the expected log likelihood, but by invoking the optimization step only after all other parameters have converged, greatly increases the computational efficiency of the algorithm. Beyond the present auction model, this GEM algorithm could be applied to any setting where the parameters to be estimated can be partitioned into a set for which analytical solutions are available and a set for which numerical methods are required to compute. The technical details for implementing the algorithm are discussed in detail in Appendix 1.A.4.

1.6.5 Monte Carlo Simulations

Two Monte Carlo Simulation exercises evaluate the performance of maximum likelihood estimation in the level-k model. The first exercise draws valuations from a log-normal distribution with mean parameter, $\mu = 0$ and standard error parameter $\sigma = 0.5$ truncated at the 99.99th percentile to ensure bounded support. Estimating the level-k auction model uses a correctly-specified parametric model having observed 60, 120, or 300 bids from simulated bidders competing in auctions with $N \in \{3, 4, 5, 6, 12\}$ for total sample sizes of 300, 600, and 1,500 bids. The second exercise draws valuations from a Legendre Sieve distribution with unit support and parameter vector $\theta = (-0.25, -0.05)$ to test the sieve estimator in a properly specified model with $N \in \{2, 3, 4\}$ for a total sample of 600 bids. Both simulations implement a model with three behavioral types: Equilibrium, Random Level-1, and Random Level-2 representing 20%, 60%, and 20% of the population, respectively.

Table 1 presents the estimation results from a set of 100 simulations under the log-normal specification. The estimator retains consistency with the MSE diminishing at roughly the expected rate as the number of observations increases. However, it is worth noting that, while estimates are very precise relating to the parameters governing the distribution over valuations, the estimates for the distribution over types are still quite noisy, maintaining a standard deviation around 10% even with 1,500 observations. The convergence properties of the Legendre sieve estimator are quite similar.

These simulations indicate that, while estimates the distribution over valuations have reasonable accuracy and precision, it is difficult to get statistically significant findings differentiating the distribution for models of bidders that are more sophisticated than the $L1_R$ bidder type. The weak estimation highlights the importance of allowing for partial identification for the model, motivating the focus in the next section on developing applied mechanism design strategies that are robust to unidentified distributions over bidder-types. Also, the weak separation among higher-order types indicates the most relevant empirical distinction is between higher-order bidding behavior where individuals account for others' bid shading and lower-order bidding behavior responding to uninformative model of others.

1.7 Expected Revenue & Optimal Reserve Price

One of the primary applications in analyzing auction data is to facilitate mechanism design decisions such as optimally setting the reserve price or choosing between a first and second price auction. Crawford, Kugler, Neeman, and Pauzner (2009) extend the Crawford and Iriberri (2007a) model to auctions with a reserve price. Their analysis illustrates the effects of behavioral agents on optimal auction design with representative examples by focusing on simple settings with two bidders when the mechanism designer knows these bidders' types. Appendix 1.A.5 develops this analysis further, characterizing expected revenues in auctions with more than two bidders under a general distribution over valuations where the composition of bidder-types in the population is unknown and may be unidentified.

The application here simply identifies the optimal reserve price in a first-price auction, which is quite narrow in scope relative to designing the revenue-maximizing mechanism. Since revenue equivalence fails in the presence of level-k bidders, Myerson (1981)'s optimal auction result does not apply so there is no reason to expect a first-price auction with reserve prices to be an optimal mechanism. To illustrate this, Crawford, Kugler, Neeman, and Pauzner (2009) present an exotic mechanism that generates greater expected revenues than is attainable in the first price auction with reserve price. Focusing on a relatively simple deviation from the original mechanism used to estimate the model provides greater confidence in the counterfactual analysis, in particular regarding bidders' response to changes in the reserve price. Implementing dramatic changes in the structure of bidding and allocation rules could affect players' participation decisions as well as their position in the behavioral hierarchy. In this case, counterfactual analysis would be misleading if it were based on the distribution over valuations and bidder-types compatible with a non-binding reserve price.

This section uses the data generating process for the winning bid from equation 1.17, treating the reserve price as a parameter in this distribution, to analyze the problem of setting a minimum reserve price for a risk-neutral seller that maximizes expected revenues when the number of participating bidders and the distribution over bidder-types is known. When the distribution over types is unidentified, the expected revenue at a given reserve price is partially identified, belonging to a compact, convex set, so that the optimal reserve price belongs to a compact identified set.

1.7.1 Calculating the Optimal Reserve Price

Given the analysis in Crawford, Kugler, Neeman, and Pauzner (2009) and in Appendix 1.A.5 defining the behavioral model, simply rewrite equation 1.17 with the reserve price, denoted r, as an additional parameter in the distribution for the value of the winning bid as $f_W(w; p, \theta, r)$.

The expected utility to a seller who attaches the value v_s to the object at auction is given by:

$$E[U_s|r] = v_s F_X(r;\theta)^N + \int_r^{\overline{x}} w f_W(w; p, \theta, r)\, dw \tag{1.24}$$

Then first order conditions for maximizing the seller's welfare are:

$$v_s N F_X(r;\theta)^{N-1} f_X(r;\theta) + \int_r^{\overline{x}} w \frac{\delta f_W(w; p, \theta, r)}{\delta r} dw - r f_W(r; p, \theta, r) = 0$$

which, similar to the Myerson (1981) analysis gives an optimal reserve price as the solution to a fixed point problem:

$$r = v_s \frac{N F_X(r;\theta)^{N-1} f_X(r;\theta)}{f_W(r; p, \theta, r)} + \frac{\int_r^{\overline{x}} w \frac{\delta f_W(w;p,\theta,r)}{\delta r} dw}{f_W(r; p, \theta, r)} \tag{1.25}$$

Equation 1.25 provides an implicit solution for the optimal reserve price given the distribution of bids in 1.17. However, in practice simulation methods that choose the optimal reserve price to maximize the conditional expected revenues in 1.24 are well-adapted to the problem.

1.7.2 The Optimal Reserve Price with Non-Identification

When there is no variation in the number of bidders or when that variation is endogeneous, the exclusion restrictions establishing identification no longer apply. To address reserve pricing in this setting, first assume that the seller knows the unconditional distribution over bidder-types with certainty, as in the previous section. This exercise can be repeated for any distribution over bidder-types with the optimal reserve prices for every such distribution characterizing the identified set for the optimal reserve price given the available information.

To establish compactness of the identified set for the optimal reserve price, note that the right hand side of equation 1.25 is continuous in changes to the distribution over bidder-types and bounded away from zero due to the regularity conditions that ensure continuously differentiable bidding strategies for all bidder-types. First, as can be verified by observing that all elements in equation 1.17 are continuous polynomials in p_k, $f_W(r;\theta,r)$ is continuous in the distribution over bidder-types. Similarly, as is shown in Crawford, Kugler, Neeman, and Pauzner (2009), each of the level-k bidder-types behavioral strategies are continuous in the reserve price. As such, $\frac{\delta f_W(w;\theta,r)}{\delta r}$ is continuous in p_k. Finally, searching over the set of bidder-types reveals the maximum and minimum reserve price that characterize the identified set.

Theorem 1.6 (Optimal Reserve Price with Non-Identification)
In the level-k auction model where the distribution over bidder-types is not identified, the optimal reserve price characterized by equation 1.25 belongs to a partially identified compact set.

1.7.3 The Ambiguity Robust Optimal Reserve Price

The selection of an optimal reserve price in a partially- or incompletely-identified model can be viewed as an exercise in decision making under ambiguity. The axiomatic choice framework introduced by Gilboa and Schmeidler (1989) rationalizes a robust decision rule that maximizes expected utility generated from the state in the identified set that minimizes expected utility conditional on the chosen action.[9]

$$r_A^* = \arg\max_r \left\{ \min_p v_s F_W\left(r; p, \theta, r\right) + \int_r^{\overline{x}} w f_W\left(w; p, \theta, r\right) dw \right\} \tag{1.26}$$

These preferences are readily applied to estimating the ambiguity-robust optimal reserve price in incompletely-identified models. For a given reserve price, the distribution over bidder-types generating the minimum expected utility to the seller will be degenerate, placing all mass on the single bidder-type that minimizes the seller's expected revenue at that reserve price. As illustrated in Crawford, Kugler, Neeman, and Pauzner (2009), while the bidder-type generating the minimum expected revenue may depend on the distribution over valuations, this regularity feature facilitates computing the min-max utility. Further, the optimal reserve price in heterogeneous populations corresponds to the optimal reserve price for the homogeneous population that minimizes the expected revenue, as in the next theorem.

Theorem 1.7 (Ambiguity Robust Optimal Reserve Price)
Suppose the distribution over bidder-types is not identified in the level-k auction model. The unique ambiguity-robust optimal reserve price maximizes the seller's expected revenue when bidders are drawn from a homogeneous population of the bidder-type $\underline{k} \in \mathcal{K}$ *that minimizes the seller's revenue.*

While beyond the scope of the current exercise, this result can be readily extended to other incompletely identified models of auction behavior. In particular, a direct corollary of theorem 1.7 applies to setting optimal reserve prices using the risk-neutral model of bidding behavior. For example, under the HARA utility specification with non-risk loving preferences, the risk-neutral model gives rise to the most aggressive bid-shading in equilibrium, effectively minimizing the expected

[9] Axiomatic treatments of decision in the presence of ambiguity date to Savage (1954), coming into stark focus with the Ellsberg (1961) paradox. Recent advances, in this area, including Klibanoff et al. (2005), Maccheroni, Marinacci, and Rustichini (2006), and Klibanoff, Marinacci, and Mukerji (2009) have extended Gilboa and Schmeidler (1989)'s results to allow for smooth preferences over ambiguity. These more complicated preferences are not as easily adapted to the current exercise, which is greatly simplified by the corner solutions imposed within the objective function by minmax preferences.

revenue over possible bidder utility-types. This result indicates that optimal reserve pricing strategies based on risk neutral-bidding behavior satisfy a robustness property even though the risk-preferences may be misspecified in the population.

1.8 Optimal Reserve Pricing in USFS Timber Auctions

With well established publicly available data, timber auctions sponsored by the US Forestry Service have received a great deal of attention in the literature on empirical methods for optimal mechanism design. This setting, then, provides an ideal environment to compare the mechanism design implications of a level-k behavioral model for IPV first price auctions with equilibrium results.

I use data provided by Philip Haile that has been used extensively in empirical studies of auctions. Early studies looking at this data include Baldwin, Marshall, and Richard (1997), who provide a detailed institutional analysis of the auctions in testing for collusion among bidders. More recent studies in empirical industrial organization by Athey and Levin (2001), Athey, Levin, and Seira (2008), and Haile and Tamer (2003) have focused on mechanism design issues in USFS timber auctions. Haile (2001) looks at the role of resale in affecting valuations for timber auctions and Haile, Hong, and Shum (2003) use USFS data to test for common value components in bidder valuations. Campo, Guerre, Perrigne, and Vuong (2000) and Lu and Perrigne (2008) use USFS auction data to characterize risk aversion within the bidding population.

The data treatment is based on the results from Haile, Hong, and Shum (2003), whose findings support the IPV model for sealed-bid timber auctions of scaled sale contracts. In these contracts, logging companies pay a price for timber harvesting rights based on the actual timber harvested, greatly reducing common- and affiliated-value components in determining the individual firm's valuation. The sample focuses on sealed bid sales from 1982-1996 that had between two and four bidders, excluding salvage sales, tracts set aside for sale to small businesses, and auctions that had more than 4 bidders. Bids from auctions in the highest and lowest 1% quantiles of appraised values are trimmed from the sample, though this had little impact on the results. For completeness, Table 1.2 presents summary statistics characterizing the entire sample of bids, though only 744 of these observations are selected after trimming.

Figure 1.7 presents the estimated distributions over valuations under homogeneous bidder-type specifications. Panels A and B present the estimation results for the level-k bidder-types estimated using several different specifications for the polynomial order of the SNP-ML Estimator. The patterns across estimation models are largely as expected. The $\mathrm{L1}_R$ bidder-type's distribution over valuations in Panel A is scaled and left-shifted relative to the distribution over bids, with the SNP-ML estimator converging quickly to the kernel based estimator based on 1.7. The Equilibrium bidder-type's distribution over valuations in Panel C is estimated using the Guerre et al. (2000) estimator, with a substantially fatter tail than either of the distributions recovered from a model that assumes Level-k bidder-types. This fatter tail is consistent with the implication that less sophisticated bidders would be over-bidding relative to the equilibrium bidder-types.

The $\mathrm{L2}_R$ bidder-type's distribution displays an interesting feature related to the potential for overfitting the model that is not be captured directly through the likelihood ratio. In particular, the $\mathrm{L2}_R$ bidder model fits the data by making the derivative of the bidding function as small as possible near the mode of the distribution over valuations, creating a spike in the likelihood. This spike is mitigated by the assumption that bidding behavior is strictly monotonic, as formalized in Lemma 1.1, though explicitly incorporating this restriction into the estimation procedure is not entirely trivial. Henderson et al. (2008) look at ways to enforce monotonicity in the estimation process for auctions under the equilibrium behavioral model based on kernel density estimation methods, though these are not immediately applicable to the sieve estimation strategy. Another approach would be to implement a constrained sieve estimator or to use a penalized likelihood criterion function as the basis for estimation.

The last panel plots the identified set implied by the observed distribution over bids. This panel illustrates the result that, even when the distribution over valuations is not identified, the bounds on the distribution over valuations are quite informative. This finding is consistent with simulation evidence that the ability to empirically separate types is relatively weak while not necessarily hindering meaningful inference on the primitives of the auction model.

Turning towards the optimal reserve price begins by briefly addressing the specification of the truncation algorightm for the SNP-ML. Model Selection Statistics for the SNP-ML estimator of the distribution over bids in homogeneous populations are presented in Table 1.3. Likelihood statistics generally favor the Level-2 Random bidder model with a very flexible distribution over bids. However, the visual evidence of over-fitting for this model is too great to ignore, so the analysis proceeds with the distribution over valuations estimated from the Level-2 Random bidder-type truncating to a 5th order Legendre polynomial.

Figure 1.8 presents counterfactual evidence on the effect of changing the reserve price for a population of $N = 4$ bidders. First, note that the bidder bidding strategies separate as expected, with more aggressive bid shading by the higher-level types. The Level-1 Random bidder type is particularly insensitive to the reserve price, which is consistent with that bidder-type's insensitivity to the distribution over valuations in choosing their bid shade. The Level-2 and Equilibrium bidder-types show significant strategic responses to the reserve price, as they no longer account for bid shading behavior below the reserve price.

The effect of the reserve price on the revenue from the auction is depicted in Figure 1.9, which plots the expected revenue from the auction at various levels of the reserve price. As is evident in Figure 1.9, the identified set for the optimal reserve price is quite large, driven mainly by the Level-1 Random bidder type's lack of sensitivity to the reserve price. Indeed, the optimal reserve price in an auction with a population entirely formed of Level-1 Random bidder-types would be equal to the seller's own valuation for the good. This feature helps to rationalize the fact that observed reserve prices in practice seem "too low" and non-binding in empirical analysis, as the cautionary reserve price could reflect skepticism on the part of the seller that higher-order bidder types dominate the population.

As discussed in the previous section, the homogeneous population consisting entirely of bidder-types who follow the most aggressive bid shading strategy also minimize expected revenues. To this ends, the ambiguity robust optimal reserve price remains the optimal reserve price estimated under the equilibrium bidding model. Worth noting, however, is that the optimal reserve price is quite far out in the tail, resulting in a nearly 70% chance that the auction will close without a buyer. For this reason, previous researchers have argued that the non-binding reserve price is likely due to non-revenue motives related to forest management and resource development. The analysis here implies uncertainty regarding the bidders' strategic response to the reserve price also rationalizes this policy.

1.9 Evaluating Bidding Behavior in the Field

Within a controlled experimental setting, the questions of identification and estimation are moot given that the distribution over valuations is controlled by design and does not need to be inferred from the bidding behavior. However, incorporating behavioral models as the basis of the econometric model not only insures against model-misspecification, but also provides a mechanism for evaluating the external validity of behavioral patterns observed in the laboratory.

As Levitt and List (2007) argue, field studies provide valuable external validation for experimental findings that characterize the economic impact of non-equilibrium behavior outside of highly controlled settings. While existing empirical studies fail to reject the BNE behavioral model, these results lack power against unstructured alternatives and can be enhanced by directing power toward behavioral alternatives. To test the equilibrium behavioral hypothesis using auctions in the field, I introduce a novel set of data representing bidding behavior in sealed bid, first price auctions for vintage computers held over the internet by the Alameda County Computer Resource Center (ACCRC) and administered by Sellam Ismail. These auctions are presently ongoing and so I hope to expand the data set beyond the very small sample of 186 bids for 55 items from 53 different bidders. As such, it is best to interpret this exercise as a pilot for a more involved field experiment evaluating individual sophistication in the field.

By tracking individual bidders across auctions, the structure of the ACCRC data greatly aids the strength of identification. The most interesting problem to study here is the behavioral patterns within this population relative to other economic settings such as timber auctions. While many of the bidders are intimately familiar with the goods at auction, they are not professionals with a great deal of experience in auctions. As such, bidder sophistication in this setting could differ substantially from the timber auction setting where bids are chosen by professionals with experience both with the auction and other bidders.

As with collectible antiques, there may be some common value component to the bidding based on the potential resale value of the vintage computers. To control for this common value component, the analysis incorporates the estimated market values calculated prior to the auction and made publicly available to all bidders by Mr. Ismail. As such, while a more robust treatment of the issue is warranted, individual valuations can be defended as independent conditional on this public information.

Table 1.4 reports summary statistics for the ACCRC bidding data sample. Fitting the sample with the Legendre sieve estimator results in recovering the distribution over valuations displayed in Figure 1.10 for the 3rd through 7th order sieve polynomial. Table 1.5 reports the estimated distribution over bidder-types along with the expected log likelihood value and the Bayesian Information Criteria (BIC) for each of these distributional specifications. These pilot results are somewhat surprising, especially considering the relatively inexperienced bidders in the auctions. Under each specification, the level-k behavior model shows the equilibrium bidder-types represent almost the entire population of bidders, with a simple BIC test selecting for the polynomial order of fit selects a 6th order polynomial.

This pilot study could be extended into a broader field experiment testing behavior in auctions in the spirit of Lucking-Reiley (1999) or List and Lucking-Reiley (2000) using online auctions. Given an appropriate market, such as that for baseball cards or other collectibles when a secondary market provides information to control for potential resale values, one could test a variety of auction formats and quickly accumulate a large amount of data on how individuals in the field play auctions. With a slightly more structured market, this setting could controlling for challenges to identification such as endogenous participation.

1.10 Conclusion

This paper proposes a structural econometric model for analyzing auction data when bidder behavior is governed by a level-k behavioral model, establishing identification conditions for the model and developing a nonparametric consistent estimation strategy. I apply the model to field studies evaluating the level of sophistication by bidders in various settings and to the applied mechanism design problem of finding the optimal reserve price in first price auctions with heterogeneous non-equilibrium behavior. These results underscore the degree to which behavioral misspecification can affect counterfactual analysis.

A possible extension of these results could address heterogeneity in both strategic beliefs and risk aversion by combining information from multiple auction mechanisms. For example, Lu and Perrigne (2008) leverage a second-price auction where an individual's decision is free of strategic and risk considerations to identify the distribution over valuations and a first-price auction to identify the bidder's utility functions. Another approach might be to analyze a model similar to Campo, Guerre, Perrigne, and Vuong (2000) and exploiting random variation in the reserve price to pin down sufficient quantiles of the distribution over bids to identify the distribution over bidder-types. Li (2005) and Li and Perrigne (2003) consider the identification problem with random reserve prices, though since the reserve prices are hidden, Li (2005) and Li and Perrigne (2003) show the uncertainty introduced to the bidding problem complicates identification rather than generating additional information on bidding characteristics in the population.

The mechanics underlying the SNP-ML consistency proof invoke distributional epiconvergence, which can generate asymptotic distributional results including central limit theorems. Delving deeper into the weak epiconvergence results may characterize the distribution for estimated parameters and statistics to test the behavioral models. Further, these results could hold under relaxed assumptions, in particular relating to assumption 1.6. Further, these results would not rely on discretizing the support for auction-specific covariates, providing general asymptotic distributional results that would also be useful in addressing finite-dimensional problems as well.

Adopting non-equilibrium behavioral economic model for structural econometric analysis could yield interesting insights in other strategic environments, such as the estimation of static and dynamic games often studied in empirical industrial organization. One such application could develop a level-k behavioral econometric model for static entry games in markets, such as those pioneered by Bresnahan and Reiss (1991) and Berry (1992). Ciliberto and Tamer (Forthcoming 2009) analyze this problem for airlines when the equilibrium is partially identified, finding that the equilibrium prediction underestimated coordination among airlines in reaction to a change in regulatory policy. Several authors, including Rapoport, Seale, and Winter (2002) and Camerer, Ho, and Chong (2004) show that players in the lab often achieve better ex-post coordination than equilibrium predicts and that this coordination is consistent with a cognitive hierarchy model. Using the Aradillas-Lopez and Tamer (2008) approach to estimating games based on rationalizability assumptions, akin to the analysis in Collard-Wexler (2008) or in a dynamic context following Aguirreagabiria and Magesan (2009) provide two interesting potential methods for relaxing the equilibrium assumption in this context.

Tables and Figures

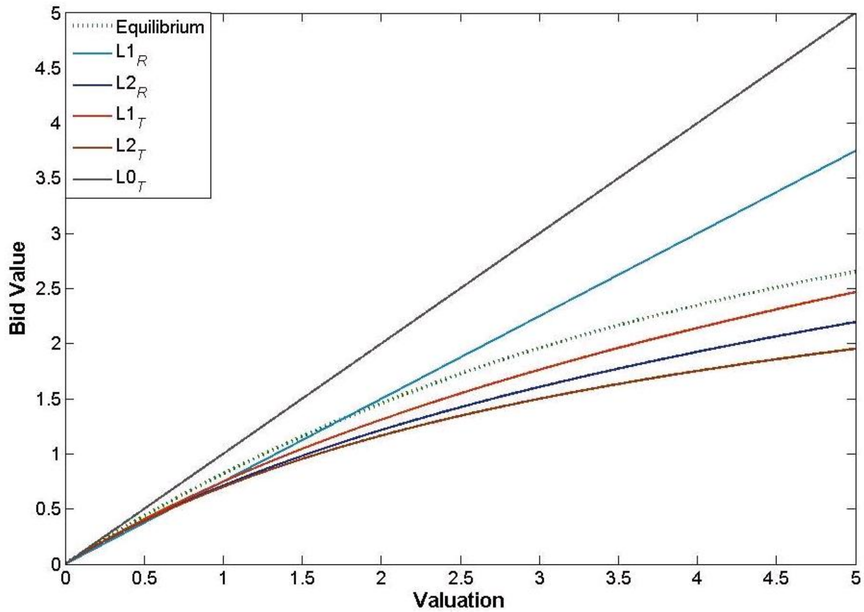

Figure 1.1: Bidding Functions for Level-k Behavioral Types
Each bidder-type corresponds to a unique, monotonic, continuous bidding strategy. Further, each pair of bidding strategies satisfies a single-crossing property that allows the econometrician to separate them.

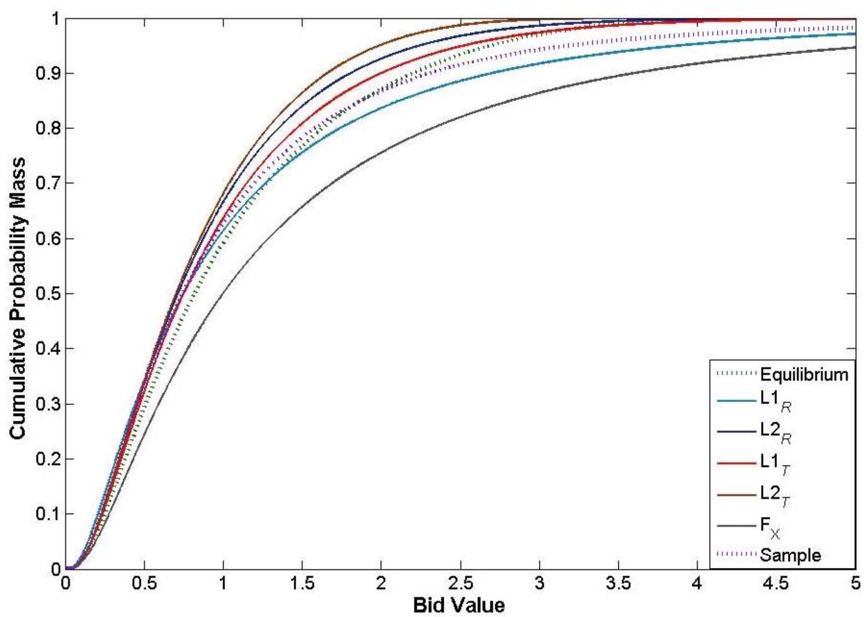

Figure 1.2: Distribution over Bids for Level-k Behavioral Types
Each bidder-type corresponds to a unique distribution over bids. The Sample distribution over bids is the mixture of these distributions that is observed by the econometrician. Note that the tail of the Sample distribution is substantially fatter than the tail for the distribution associated with the equilibrium bidder-type, so we'd expect estimation based on the equilibrium model to recover a distribution over valuations with substantially fatter tails.

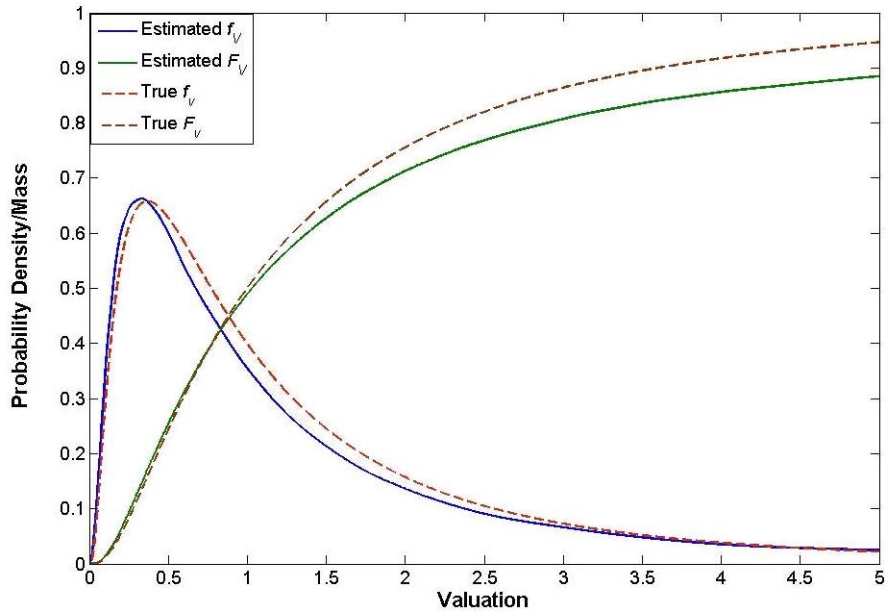

Figure 1.3: Incorrect Inference Caused by Behavioral Misspecification
Assuming an equilibrium behavioral model in the presence of behavioral bidders can lead to substantial errors in inference. The estimated distribution represented by solid lines has a substantially fatter tail than the true distribution represented by the dashed line.

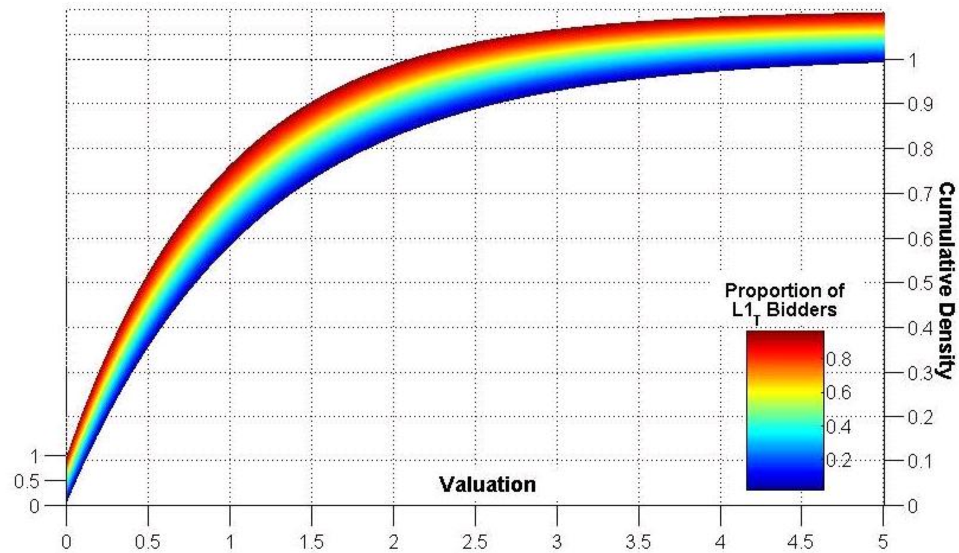

Figure 1.4: Identified Set when $\mathcal{K} = \{L0_T, L1_R\}$, $N = 5$
The identified set when $N = 5$ bidders with an unknown distribution over Level-0 Truthful and Level-1 Random bidder-types includes a unique distribution over valuations for any mixture of the bidder-types.

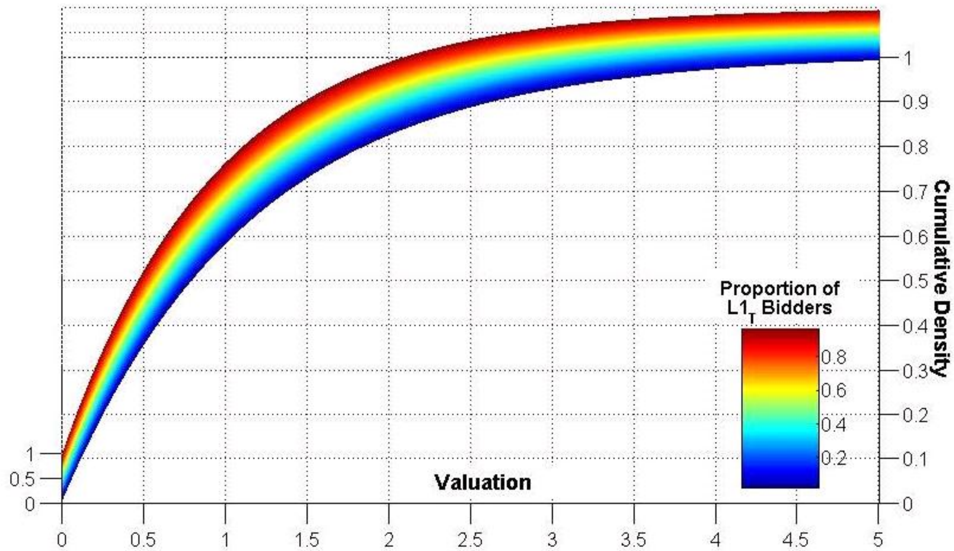

Figure 1.5: Identified Set when $\mathcal{K} = \{L0_T, L1_R\}$, $N = 20$
The identified set when $N = 20$ bidders with an unknown distribution over Level-0 Truthful and Level-1 Random bidder-types includes a unique distribution over valuations for any mixture of the bidder-types. Note that the identified set has shifted from the $N = 5$ case and, in particular, every estimated distribution except the true distribution over valuations, corresponding to the true distribution over bidder-types, has shifted.

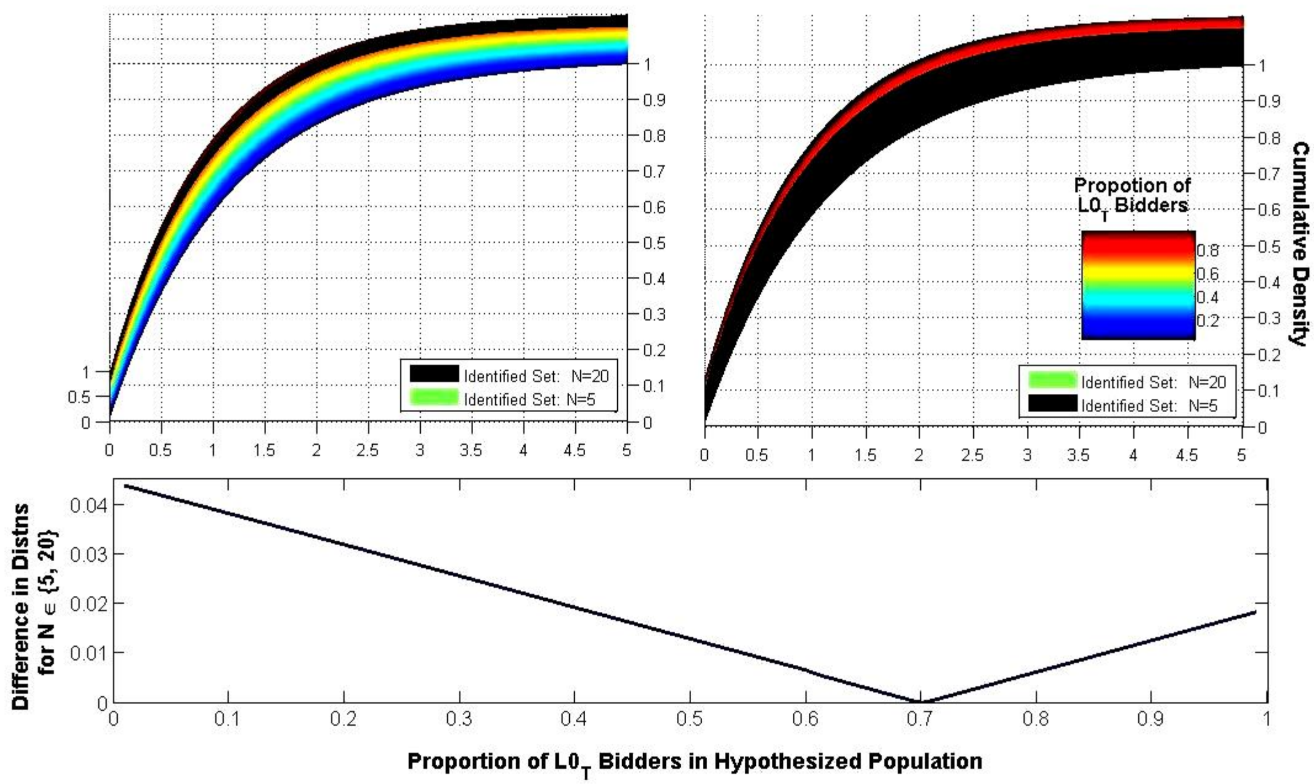

Figure 1.6: Identification from Intersecting Identified Sets for $N \in \{5, 20\}$
The top two panels show that when we plot the identified sets corresponding $N = 5$ bidders and $N = 20$ bidders, the intersection is unique. The bottom panel shows that this intersection (where the sup-norm of the difference between the distributions is zero) identifies the true distribution over bidder-types in the population.

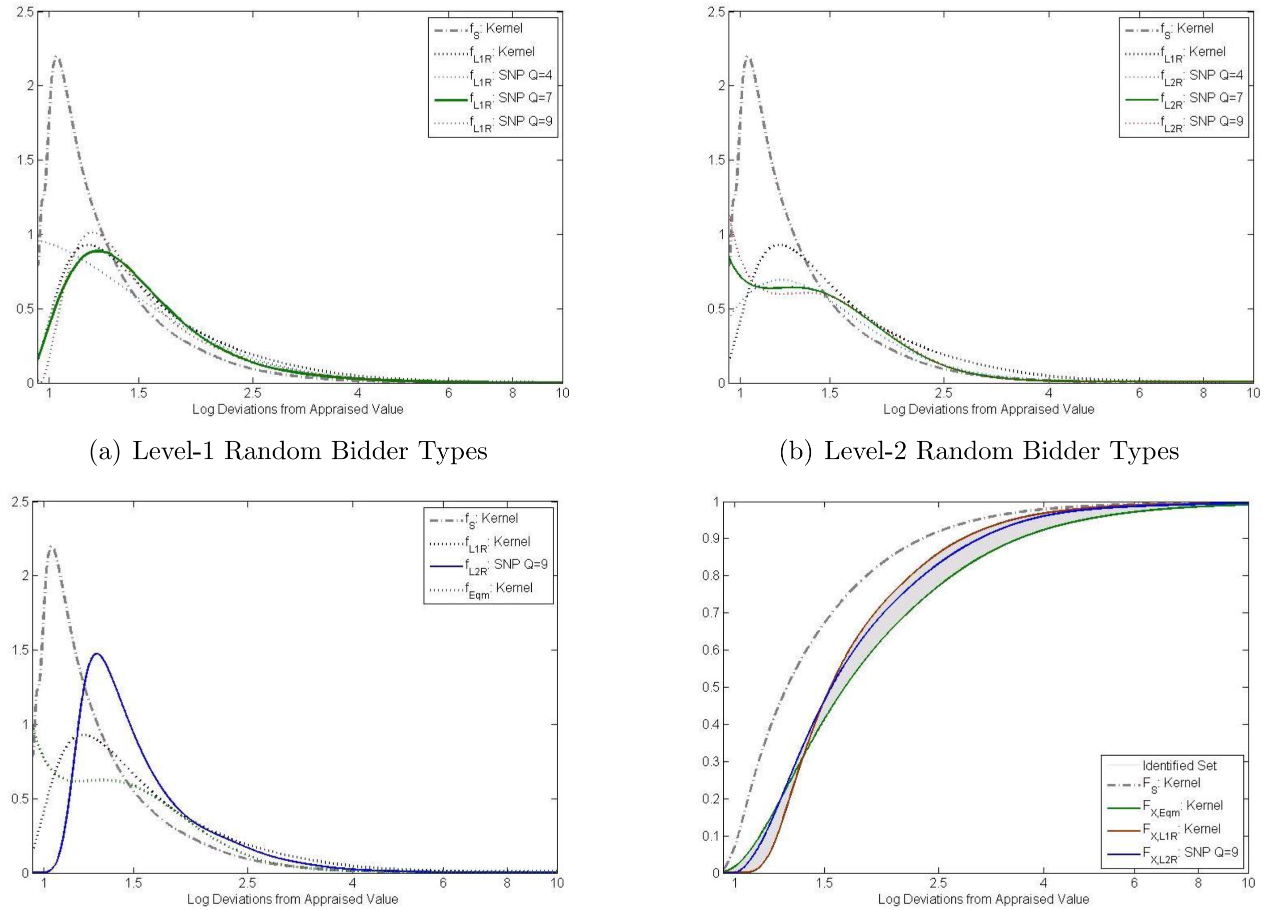

(a) Level-1 Random Bidder Types

(b) Level-2 Random Bidder Types

(c) Equilibrium Bidder Types

(d) Identified Set of Distribution over Valuations

Figure 1.7: Distributions for Valuations in USFS Timber Auctions

Panels (A) - (C) represent the distribution over valuations implied by observed bidding behavior in USFS timber auctions when the population is comprised entirely of level-1 random, level-2 random, or Equilibrum bidder-types, respectively. Panel D illustrates the identified set of cumulative densities for the true distribution over bids.

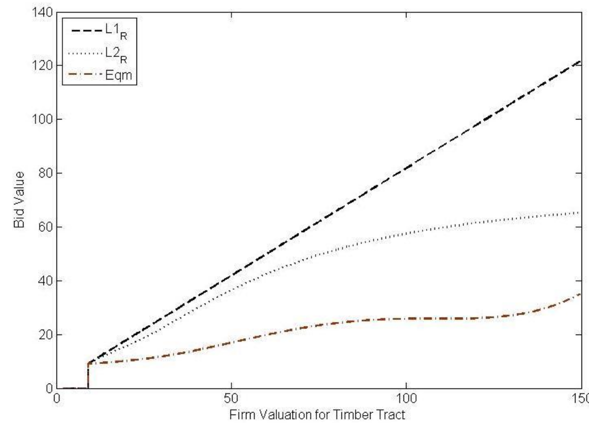

(a) All Bidder-type Strategies

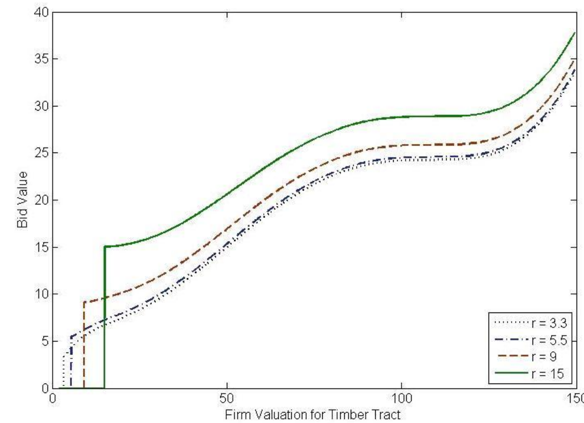

(b) Equilibrium Bidder Types

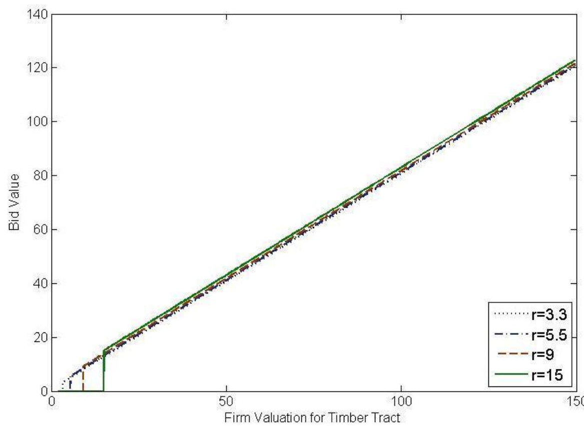

(c) Level-1 Random Bidder-Type

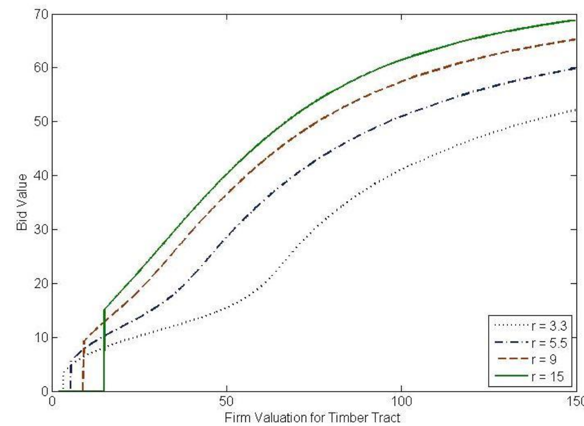

(d) Level-2 Random Bidder-Type

Figure 1.8: Implied Bid Function Responses to Reserve Prices
The bidding functions implied under alternate reservation prices for USFS timber auctions for the level-1 random bidder-type are much less sensitive than the rules followed by more sophisticated bidders. As such, the potential benefit of shifting the reservation price in a less sophisticated population is not as great as for higher order types.

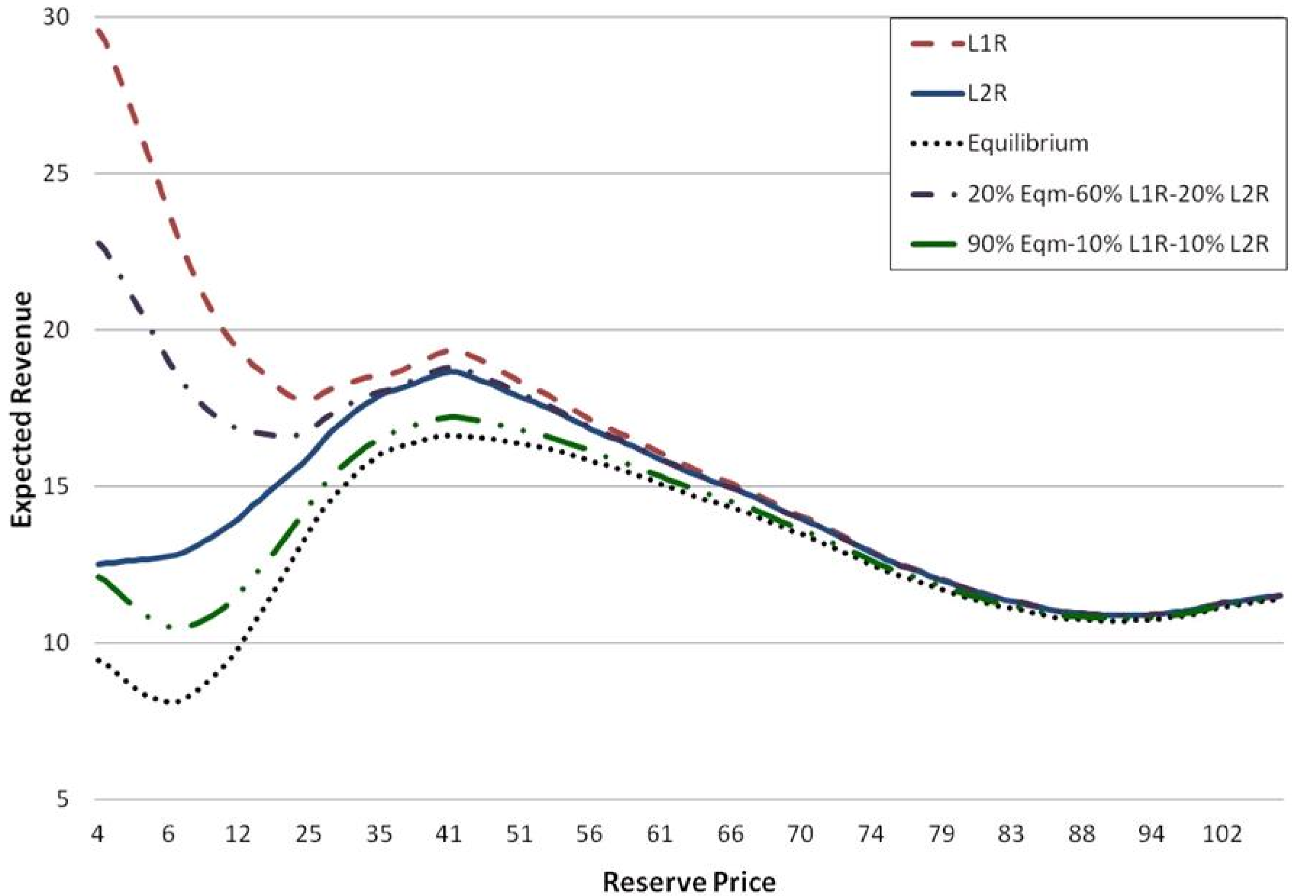

Figure 1.9: Expected Revenues for USFS Timber Auctions

This figure plots the effect of the reserve price on expected revenues in USFS Timber Auctions. When the population is dominated by the Level-1 Bidder-type, the seller maximizes revenue by setting the reserve price equal to their value for the good. When the seller faces a population with greater sophistication, the strategic benefits of the reserve price outweigh the costs from potentially losing a sale, yielding a relatively high optimal reserve price.

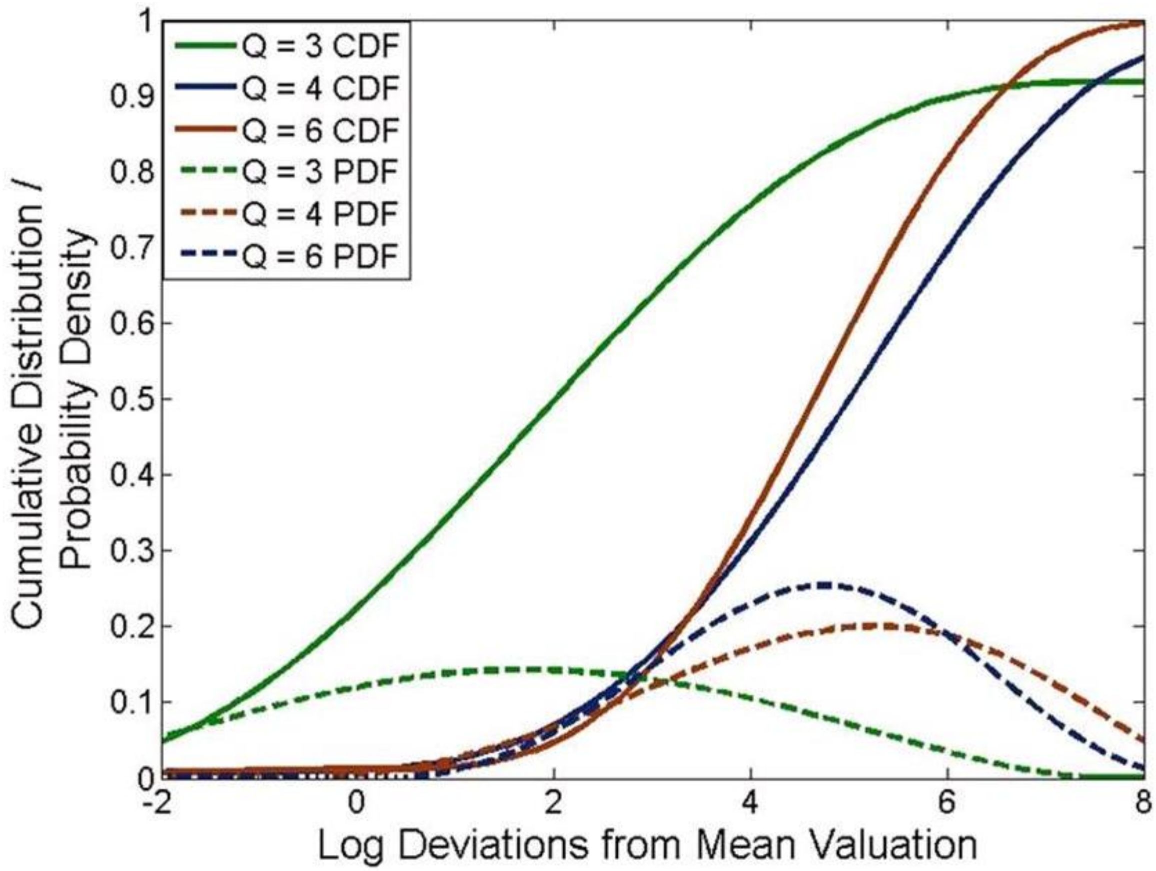

Figure 1.10: Distributions for Latent Valuations in ACCRC Auctions
This figure plots the distribution over valuations for the ACCRC vintage computer auctions in terms of log deviations from the expected valuation conditional on the estimated value provided to bidders. Though the sample size is relatively small, there is no visual effect of overfitting by the model.

Table 1.1: Monte Carlo Simulations for Parametric Estimation
This table reports the result of maximum likelihood estimation for a parametric simulation where individual valuations are drawn from a truncated Lognormal distribution and individual bids are chosen according to a randomly assigned behavioral type. These estimator results are generated from 100 simulated samples and illustrate both the consistency of the estimation strategy and the need for a robust sample size for precise estimation.

	P_{Eqm}	P_{L1_R}	P_{L2_R}	Mean	StDev
Population	0.2	0.6	0.2	0	0.5
	Log Normal Parametric Estimator				
# of Obs	Mean Square Error (*100)				
300	3.500	2.820	3.940	0.020	0.030
600	2.240	1.030	2.850	0.010	0.010
1,500	2.110	0.580	2.040	-	-
# of Obs	Standard Deviation				
300	0.185	0.168	0.197	0.013	0.018
600	0.132	0.101	0.159	0.008	0.012
1,500	0.109	0.074	0.120	0.005	0.007
	Legendre Sieve Estimator				
# of Obs	Mean Square Error (*100)				
600	2.417	0.615	2.449	0.049	0.050
	Standard Deviation				
600	0.155	0.078	0.155	0.022	0.022

37.9% of simulations resulted in corner solutions for P_{Eqm} or P_{L2_R}.

Table 1.2: Summary Statistics for USFS Timber Auction Data
This table reports summary statistics for USFS Timber Auction Data. The anlaysis uses only the auctions with 2-4 bidders, which still leaves a sample of over 744 bids for estimation.

# of Bidders	# of Auctions	# of Obs	Mean Appraisal ($000)	Mean Bid ($000)	St Dev Bid ($000)
2	143	286	886	1,215	1,690
3	75	225	810	1,276	1,512
4	68	272	1,401	2,256	3,107
5	57	285	2,480	3,861	6,912
6	26	156	1,719	2,897	3,101
7	23	161	2,518	5,112	9,325
8	6	48	701	1,188	621
9	9	81	3,245	6,585	5,907
10	13	130	4,865	11,748	19,871
Full Sample		1,655	1,892	3,481	7,780

Table 1.3: Model Selection Statistics for USFS Timber Auction Data
This table reports model selection statistics for USFS Timber Auction Data. These results are generated from observed bids in 744 auctions and provide substantial support to the hypothesis that bidders in this setting are sophisticate, though the BIC selected model appears to substantially overfit the data.

	Level 1 Random		Level 2 Random	
Polynomial Order	Log Likelihood	BIC	Log Likelihood	BIC
3	(10,324)	20,674	(10,148)	20,322
4	(10,309)	20,652	(9,979)	19,990
5	(10,298)	20,635	(9,775)	19,589
6	(10,295)	20,636	(9,774)	19,594
7	(10,291)	20,635	(9,740)	19,533
8	(10,286)	20,632	(9,660)	19,379
9	(10,284)	20,634	(9,652)	19,369
10	(10,283)	20,638	(9,629)	19,330

Table 1.4: Summary Statistics for ACCRC Auction Data
This table reports summary statistics for ACCRC Vintage Computer Auction Data. The present sample is too small for conclusive findings, so the treatment of this data is best viewed as part of a pilot for a broader field study in auction behavior.

# of Bidders	# of Auctions	# of Obs	Mean Bid	St Dev Bid
2	6	12	61.92	111.38
3	11	33	30.78	29.00
4	9	36	42.94	48.00
5	3	15	35.99	19.57
12	1	12	135.33	125.65
14	1	14	254.29	196.08
16	1	16	263.67	265.04
Full Sample		138	95.99	150.56

Table 1.5: Distribution over Bidder Types for ACCRC Auction Data
This table reports estimated population characteristics and model selection statistics for ACCRC Vintage Computer Auction Data. The findings support equilibrium bidding behavior, though these come with little confidence given the limited sample size.

Polynomial Order	P_{Eqm}	P_{L1_R}	P_{L2_R}	Likelihood	BIC
Q = 3	100.0%	0.0%	0.0%	420.52	6.19
Q = 4	90.5%	0.0%	9.5%	305.21	4.54
Q = 5	90.5%	0.0%	9.5%	287.9	4.32
Q = 6	91.0%	0.0%	9.0%	138.24	2.17
Q = 7	91.0%	0.0%	9.0%	139.03	2.21

Appendix 1.A.1: Proofs

Proof of Theorem 1.1

As in the case of the Truthful Level 1 bidder type, the identification proof begins by characterizing the identified set as a set of distributions consistent with the equality:

$$F_X\left(\sigma^{-1}(x)\right) = F_S(x) = F^*\left(\sigma_*^{-1}(x)\right) \tag{1.A.1}$$

The key here is assumption 4, which ensures an individual's bid is equal to their valuation minus a non-negative, continuously differentiable bid shade that is zero for valuations arbitrarily close to zero. These properties establish two contradictions to complete the proof.

First, suppose that $\epsilon_1 \equiv \inf\{x : F_X(x) \neq F^*(x)\} > 0$ and go on to define $\epsilon_2 \equiv \inf\{x > \epsilon_1 : F_X(x) = F^*(x)\} > \epsilon_1$ so that for $y \in [0, \epsilon_1)$, $F_X(y) = F^*(y)$ and, as such, the inverse bidding functions are identical to one another in this region, i.e., $\sigma^{-1}(y) = \sigma_*^{-1}(y)$. Note that $\sigma^{-1}(y)$ is always greater than y, continuous, and strictly increasing, so there is some $\tilde{y} < \epsilon$ with $\sigma^{-1}(\tilde{y}) \in (\epsilon_1, \epsilon_2)$. Then, $F_X(\sigma^{-1}(\tilde{y})) \neq F^*(\sigma_*^{-1}(\tilde{y}))$, contradicting 1.A.1. As such, any candidate distribution satisfying the condition 1.A.1 must differ from the true distribution of valuations starting at the origin.

Now, suppose the distributions F_X and F^* diverge immediately from the origin and that $F_X(x) > F^*(x)$ for $x \in [0, \epsilon)$ where $\epsilon \equiv \sup\{x : F_X(x) > F^*(x)\}$. In this case, the condition in equation 1.A.1 demands that $\sigma_k^{-1}(x) < \sigma_{k,*}^{-1}(x)$. However, there must come a point in $[0, \epsilon]$ where the distribution F^* begins "catching up" with F_X, i.e., where $f_X(x) < f^*(x)$. Further, by iterating down the hierarchy of bidder-types, these dual inequalities imply the Jacobian terms also satisfy: $\frac{d\sigma_{k-1}^{-1}(s)}{ds}|_{s=\sigma_k(x)} < \frac{d\sigma_{k-1,*}^{-1}(s)}{ds}|_{s=\sigma_{k,*}(x)}$. The three inequalities combined imply:

$$\frac{F_X(x)}{f_X(x)\,\frac{d\sigma_{k-1}^{-1}(s)}{ds}|_{s=\sigma_k(x)}} > \frac{F^*(x)}{f^*(x)\,\frac{d\sigma_{k-1,*}^{-1}(s)}{ds}|_{s=\sigma_{k,*}(x)}} \tag{1.A.2}$$

This result implies the bid shade under the alternative distribution, F^* is *greater* than the bid shade under the true distribution, contradicting the requirement that $\sigma_k^{-1}(x) < \sigma_{k,*}^{-1}(x)$ and proving the result.

Proof of Theorem 1.3

First, suppose $K = 2$ with known bidding strategies $\sigma_1(x)$ and $\sigma_2(x)$, the first step is to separate the mixture distribution of bids into the distribution over bids for homogeneous populations and use these components to recover the distribution over valuations. Here, the mixture distribution over bids can be written as:

$$F_{S,N}(x) = \alpha_1 F_X(\sigma_1(x)) + (1 - \alpha_1) F_X(\sigma_2(x)) \tag{1.A.3}$$

Defining $\sigma_{2\to1}(x) \equiv \sigma_1^{-1}[\sigma_2(x)]$, as the signal bidder-type 1 would need to observe to choose the same bid as bidder-type 2, rewrite 1.A.3 so as to focus on the distribution of valuations by:

$$\begin{aligned} F_{S,N}\left(\sigma_1^{-1}(x)\right) =& \alpha_1 F_X(x) + (1-\alpha_1) F_X\left(\sigma_1^{-1}[\sigma_2(x)]\right) \\ =& \alpha_1 F_X(x) + (1-\alpha_1) F_X(\sigma_{2\to1}(x)) \end{aligned}$$

This expression recovers the distribution over valuations as:

$$F_X(x) = \frac{1}{\alpha_1} F_{S,N}\left(\sigma_1^{-1}(x)\right) - \frac{1-\alpha_1}{\alpha_1} F_X(\sigma_{2\to1}(x)) \tag{1.A.4}$$

Assume (wlog) that $\alpha_1 > \frac{1}{2}$, and use the equation 1.A.4 as the basis for iteratively defining the distribution over valuations as a function of the bidding distributions and strategies, since:

$$\begin{aligned} F_X(\sigma_{2\to1}(x)) =& \frac{1}{\alpha_1} F_{S,N}\left(\sigma_1^{-1}(\sigma_{2\to1}(x))\right) - \frac{1-\alpha_1}{\alpha_1} F_X(\sigma_{2\to1}(\sigma_{2\to1}(x))) \\ \equiv& \frac{1}{\alpha_1} F_{S,N}\left(\sigma_1^{-1}(\sigma_{2\to1}(x))\right) - \frac{1-\alpha_1}{\alpha_1} F_X\left(\sigma_{2\to1}^{(2)}(x)\right) \end{aligned}$$

Then write the distribution over valuations as the infinite sum:

$$\begin{aligned}F_X(x) =& \frac{1}{\alpha_1} F_{S,N}\left(\sigma_1^{-1}(x)\right) - \frac{1-\alpha_1}{\alpha_1^2} F_{S,N}\left(\sigma_1^{-1}\left(\sigma_{2\to 1}(x)\right)\right) \\ &+ \frac{(1-\alpha_1)^2}{\alpha_1^2} F_X\left(\sigma_{2\to 1}^{(2)}(x)\right) \\ =& \frac{1}{\alpha_1} F_{S,N}\left(\sigma_1^{-1}(x)\right) + \sum_{i=1}^{\infty} (-1)^i \frac{(1-\alpha_1)^i}{\alpha_1^{i+1}} F_{S,N}\left(\sigma_1^{-1}\left(\sigma_{2\to 1}^{(i)}\right)\right) \end{aligned} \tag{1.A.5}$$

Since $\sum_{t=1}^{T} \frac{(1-\alpha_1)^t}{\alpha_1^{t+1}} \to_{T\to\infty} C < \infty$ and $0 \leq F_{S,N}\left(\sigma_1^{-1}\left(\sigma_{2\to 1}^{(i)}\right)\right) \leq 1$, This last sum converges. Extending the argument to a general number of bidder-types is straightforward (though it requires somewhat cumbersome notation) when there is a dominant bidder-type, with the only challenge being to prove that the sum in equation 1.A.5 converges. When there is not a dominant bidder-type, one can be constructed as a mixed-strategy of $K-1$ bidder-types' level-k strategies and the argument proceeds inductively.

Appendix 1.A.2: Legendre Sieve Space

This appendix provides background on the Legendre Polynomial sieve and basic results due to Bierens (2006) and Bierens and Song (2007) establishing the sieve space as a compact metric space. Note that computation using these distributions is numerically challenging that is greatly aided by some clever algorithms described in Bierens (2006). These techniques are not reproduced here but readers are referred to Bierens (2006) for implementation guidance when using these procedures.

Definition 1.3 (Legendre Polynomials) *Legendre polynomials $\rho_n(x)$ with order*
$n \geq 2$ are defined recursively by the formula:

$$\rho_n(u) = \frac{\sqrt{2n-1}\sqrt{2n+1}}{n}(2u-1)\rho_{n-1}(u) - \frac{(n-1)\sqrt{2n+1}}{n\sqrt{2n-3}}\rho_{n-2}(u)$$

with $\rho_0(u) = 1, \rho_1(u) = \sqrt{3}(2u-1)$

To adapt the Legendre polynomials to density estimation, Bierens (2006) shows that any density function $h(u)$ on $[0,1]$ can be represented as:

$$h(u) = \frac{\left(1 + \sum_{j=1}^{\infty} \delta_j \rho_j(u)\right)^2}{1 + \sum_{j=1}^{\infty} \delta_k^2}, \text{ where, } \sum_{j=1}^{\infty} \delta_k^2 < \infty \tag{1.A.6}$$

While the unit-interval support addresses measurability issues regarding the arguments of equation 1.A.6, some additional constraints are needed to ensure the parameters δ_j are well-behaved to ensure compactness for the space of density functions. This constraint takes the following form:

Lemma 1.4 (Legendre Polynomial Sieve Space (Bierens & Song)) *Let $\mathcal{D}$ be the space of density functions $h(u)$ of the form 1.A.6 where, for some a priori chosen constant $c > 0$, the parameters δ_j satisfy:*

$$|\delta_j| \leq c\left(1 + \sqrt{j}\ln j\right)^{-1}, j = 1, 2, 3, \ldots$$

Then with the L^1 metric, $\mathcal{D}$ is a compact metric space. Also, letting $G(v)$ and $g(v)$ be as in Assumption 1.4, the space

$$\mathcal{D}(G) = \{f(v) = h(G(v))\, g(v), h \in \mathcal{D}\}$$

of densities on $\left[\underline{M}, \overline{M}\right]$ with the L^1 metric is also a compact metric space. Further, the corresponding spaces of absolutely continuous distribution functions on $[0,1]$ and $\left[\underline{M}, \overline{M}\right]$, respectively,

$$\mathcal{H} = \left\{H(u) = \int_0^u h(z)\, dz, h \in \mathcal{D}\right\} \mathcal{F} = \left\{F(v) = \int_0^v f(z)\, dz, f \in \mathcal{D}(\mathcal{G})\right\}$$

with the L^1 metric are compact metric spaces.

Proof. Bierens (2006) Theorems (8) and (9) and Bierens and Song (2007), Lemmas (5), (6), and (7). ■

Appendix 1.A.3: Identification and Estimation of QRE Auctions

Kagel and Roth (1997) survey much of the early experimental evidence testing equilibrium behavioral models, presenting the stylized fact that individuals tend to over-bid relative to the equilibrium prediction in IPV auctions. To address this issue, several papers have proposed behavioral models for decision making that better fit behavior in experiments, including models based on the "Joy of Winning" (proposed by Cox, Smith, and Walker (1992))[10], quantal response equilibrium play analyzed by Goeree, Holt, and Palfrey (2002), and the cognitive hierarchy approach studied here.

A leading alternative to a cognitive-based strategic model of bidding is due to Goeree, Holt, and Palfrey (2002), who present experimental evidence that a Quantal Response Equilibrium allowing for individuals to noisily perceive payoffs in responding to equilibrium behavior fits bidding behavior observed in the lab well. In the QRE framework, individual payoffs are augmented to include a noisy error term with a logistic distribution, leading to equilibrium behavior in which individuals' bids are governed by mixed strategies, distributed so that the probability weight of choosing a given bid is proportional to a transformation of the expected utility from choosing that bid. Bajari and Hortacsu (2005) exploit this property to derive an estimation strategy for the logit QRE in empirical settings but fail to address identification beyond the attainment of a global maximum for their parametric likelihood function. I provide a more formal identification argument here, using a discrete analog to Fredholm theory commonly applied in analyzing identification of games studied by researchers in empirical industrial organization. Similar to the identification results for the cognitive hierarchy under a constant level of competition, these arguments recover a distribution over valuations consistent with a given distribution of bids for any value of the parameter governing

[10]In the absence of risk aversion, the joy of winning component for an individual's valuation is integrated with the latent valuation and, as such, not independently identifiable. This model may be identifiable with parametric risk aversion and exclusion restrictions, but given the focus on risk neutrality, it falls beyond the scope of current work.

the distribution for the error term.

The approach to analyzing the logit QRE model is similar to the strategy followed in the paper. I begin by presenting a semi-parametric incomplete identification result: conditional on the parameter governing the distribution of errors, the distribution over valuations is uniquely identified even if the econometrician observes only settings with no variation in the number of bidders. Leaving aside the "rounding" issues raised by discretizing the space of bids and valuations, I show that the QRE model is uniquely identified and testable given exogenous variation in the number of bidders. This result provides a more precise characterization of the statement by Bajari and Hortacsu (2005) that "nonparametric identification of the QRE specification may not be possible in this setting if one abandons the iid assumption and allows for enough flexibility in the distribution of the idiosyncratic shock term." The positive identification result illustrates the structure needed to address the Haile, Hortacsu, and Kosenok (2008) finding that unrestricted QRE models place no empirical restrictions on observed behavior by using an error specification consistent with the Regular QRE models proposed in Goeree, Holt, and Palfrey (2005).

The QRE model for auctions and the noisy equilibrium bidding function adopts a discretized space of bids ($s \in \mathcal{S}$) and valuations ($X \in \mathcal{V}$). In practice, bids in most auctions cannot be made for fractions of pennies so this discretization provides measurability without affecting the empirical results. Whereas the utility for a player in a BNE is given by the equation 1.1 that $E\left[U_i\left(X_i, s_1, \ldots, s_N\right)|X_i\right] = \left(X_i - s_i\right) Pr\{s_i > \max_{j \neq i} s_j\}$, in the QRE, this expected utility is perturbed by a noise term:

$$E_{QRE}\left[U_i\left(X_i, s_1, \ldots, s_N\right)|X_i\right] = \left(X_i - s_i\right) Pr\{s_i > \max_{j \neq i} s_j\} + \epsilon\left(s_i, X_i\right) \quad (1.A.7)$$

Following the analysis of Goeree, Holt, and Palfrey (2002), assume that $\epsilon\left(s_i, X_i\right)$ is independent of s_i and X_i and identically distributed according to the Type-II Extreme Value distribution with cumulative distribution function $F\left(\epsilon\right) = \exp\left(-\exp\left(-\lambda\epsilon\right)\right)$ where the mean and variance of the error term are both decreasing in the value of the parameter λ. From the Bajari & Hortascu analysis, the probability that a player i chooses a bid value of s_i conditional on the QRE distribution over bids, denoted $\sigma_i(s_i; X_i, F_S(s))$ can be stated up to a constant of proportionality as:

$$\sigma_i(s_i; X_i, F_S(s)) = \frac{\exp\left(\lambda(X_i - s_i) F_S(s_i)^{N-1}\right)}{\sum_{s \in \mathcal{S}} \exp\left(\lambda(X_i - s) F_S(s)^{N-1}\right)} \tag{1.A.8}$$

The key identifying feature that the QRE behavioral model that shares with BNE is the equilibrium requirement that individuals best respond (albeit noisily) to the empirical distribution of bids observed by the econometrician. This feature greatly aids in identification analysis as it avoids the computationally demanding exercise of calculating a fixed point that would have been required without observing the empirical distribution of bids. The unconditional probability of observing a given bid s for a representative bidder can then be computed by integrating over the possible valuations that could generate that bid (following Bajari and Hortacsu (2005)), to show that:

$$F_S(s) = \int \frac{\sum_{\tau \le s} \exp\left(\lambda(x - \tau) F_S(\tau)^{N-1}\right)}{\sum_{\tau \le x} \exp\left(\lambda(x - \tau) F_S(\tau)^{N-1}\right)} f_X(x)\, dx \tag{1.A.9}$$

To establish identification, note that this equation is a discretization of a Fredholm integral equation of the first kind, since the ratio in the integral is a probability density and, as such, also an integral kernel. While the kernel is discontinuous, these discontinuities can be made arbitrarily small through the partitioning of the bid and valuation space. As is common in identification proofs from dynamic games with incomplete information (see, for instance, Bajari, Chernozhukov, Hong, and Nekipelov (2008)), the identification result for a fixed value of λ and discretization of the bid space follows immediately by observing that Fredholm Integral Equations of the first kind have a unique solution. This observation effectively completes the proof of the incomplete identification result that parallels the result of theorem 1.3. Extending the identification result to incorporate the variation in the number of bidders is also very straightforward, as the integral kernel in 1.A.9 and, as such, the solution to the integral equation itself depends deterministically on the number of bidders and the error parameter λ. This result allows consistent estimation of the distribution over values for a fixed value of λ and two different levels of competition, N_1 and N_2 to test the hypothesis that the QRE is properly specified with that level of variance in the error term.

Theorem 1.8 (Identification of Quantal Response Equilibrium)

1. *Suppose the econometrician observes the distribution over bids in a fixed population of N bidders is $F_{S,N}(x)$, then for any fixed value of the error distribution, λ, there exists a unique distribution $F_X(x)$ generating $F_{S,N}(x)$ in a Quantal Response Equilibrium behavioral model.*

2. *Suppose the econometrician observes the distribution over bids in a variable population of $N_1, \ldots N_K$ for $K \geq 2$ bidders is $F_{S,N_k}(x)$, then the parameter governing the error distribution, λ and the distribution $F_X(x)$ generating $F_{S,N_1}(x), \ldots, F_{S,N_K}(x)$ in a Quantal Response Equilibrium behavioral model are identified. Further, if $K \geq 3$, then the QRE behavioral model and error specification impose testable overidentifying restrictions on the data.*

With the identification result now firmly established, the estimation strategy presented by Bajari and Hortacsu (2005) estimate the QRE model with exogeneous variation in the number of bidders. Absent exogeneous variation in the number of bidders participating in an auction, though, it is important to remember the lack of identification power for the error parameter. Further, the identification result is highly dependent upon the Type-II Extreme Value distribution and it is likely that, for any specific distribution over error terms, there exists a unique distribution over valuations consistent with observed bidding behavior.

Appendix 1.A.4: A Generalized EM Algorithm

Given the types generating each bid, d_{ik}, the log likelihood function 1.15 is:

$$\ln \mathcal{L}\left(\theta; \{s_i\}_{i=1}^{T}, \{d_{ik}\}_{i=1}^{T}{}_{k=1}^{K}\right) = \sum_{i=1}^{T}\sum_{k=1}^{K} d_{ik} \ln f_{S,k}(s_i;\theta) + \sum_{i=1}^{T}\sum_{k=1}^{K} d_{ik} \ln p_k$$

Though this likelihood function is not directly observable, it can be approximated by taking the expectation over the unobservable d_{ik} parameters to get an expected log likelihood:

$$E\left[\ln \mathcal{L}\left(\theta; \{s_i\}_{i=1}^{T}, \{d_{ik}\}_{i=1}^{T}\ _{k=1}^{K}\right)\right]$$
$$= \sum_{i=1}^{T}\sum_{k=1}^{K} E\left[d_{ik}\right] \ln f_{S,k}\left(s_i;\theta\right) + \sum_{i=1}^{T}\sum_{k=1}^{K} E\left[d_{ik}\right] \ln p_k$$

To compute this expected likelihood, initialize the process with an a priori guess for the distributional parameters θ_0 and the distribution over types p_0. Then estimate $\hat{z}_{ik,0} = E_0\left[d_{ik}\right]$, which is the probability that bid s_i is drawn from the distribution of bids for the kth bidder-type. This probability is a straightforward application of Bayes' rule given by a formula from the mixture-of-types models of Stahl and Wilson (1994), Stahl and Wilson (1995) and Costa-Gomes, Crawford, and Broseta (2001):

$$\hat{z}_{ik,0} = p_k\left(s_i;\theta,p\right) = \frac{p_k\left(\theta\right) f_{S,k}\left(s_i;\theta\right)}{\sum_{\kappa\in\mathcal{K}} p_\kappa\left(\theta\right) f_{S,\kappa}\left(s_i;\theta\right)} \tag{1.A.10}$$

The Expectation step in the Expectation Maximization (EM) algorithm then approximates the above log likelihood by:

$$E_0\left[\ln \mathcal{L}\left(\theta; \{s_i\}_{i=1}^{T}, \{d_{ik}\}_{i=1}^{T}\ _{k=1}^{K}\right)\right] = \sum_{i=1}^{T}\sum_{k=1}^{K} \hat{z}_{ik} \ln f_{S,k}\left(s_i;\theta\right) + \sum_{i=1}^{T}\sum_{k=1}^{K} \hat{z}_{ik} \ln p_k$$

The Maximization step in the EM algorithm then chooses the parameter vector θ_1 and distribution over types p_1 to maximize this expected log likelihood and proceeds to iterate between the Expectation and Maximization steps until these distributional estimates converge. Since the parameters governing the distribution over bids for fixed bidder-type does not depend on p, the maximization problem can be separated into two pieces. First, the updated distribution over types must be the average probability that a bidder is drawn from that type. That is:

$$p_{k,1} = \frac{1}{T}\sum_{i=1}^{T} \hat{z}_{ik,0} \tag{1.A.11}$$

$$\theta_1 = \arg\max_{\theta} \sum_{i=1}^{T}\sum_{k=1}^{K} \hat{z}_{ik} \ln f_{S,k}\left(s_i;\theta\right) \tag{1.A.12}$$

Computationally, the Expectation step and the first piece of the Maximization step in the EM algorithm is quite fast, even when numerical methods are used to compute the equilibrium bidding and inverse bidding functions and their associated derivatives. However, the Maximization step that requires generating new estimates for θ is quite cumbersome due to its role in computing equilibrium bidding functions. As such, repeated application of the EM algorithm until the algorithm converges is computationally infeasible and requires a great deal of redundant calculations. Further, the convergence for the parameter estimates occcurs much more quickly than convergence for the distribution over behavioral types.

To address this issue, I introduce a Generalized EM algorithm that proceeds as follows:

Algorithm 1 ***Generalized Expectation Maximization Algorithm***

Step 0: *Initiate model with a priori guesses for* p_0 *and* θ_0*, choose tolerance* δ*, set* $p_{0a} = p_0$.

Step 1: *Expectation Step: Use equation 1.A.10 to compute* $\hat{z}_{ik,0}$.

Step 2: *Partial Maximization Step:*

Step 2a: *Use equation 1.A.11 to compute* p_1.

Step 2b: *If* $||p_0 - p_1|| > \delta$*, set* $p_0 = p_1$ *and return to Step 1.*

Step 3: *Complete Maximization Step:*

Step 3a: *Choose* θ_1 *to maximize the expected log likelihood formula in equation 1.A.12.*

Step 3b: *Use equation 1.A.11 to compute* p_1

Step 3c: *If* $||p_{0a} - p_1|| + ||\theta_0 - \theta_1|| > \delta$*, set* $p_0 = p_1$*,* $p_{0a} = p_1$*,* $\theta_0 = \theta_1$ *and return to Step 1.*

Generalized EM (GEM) algorithms are well-known tools for addressing maximum likelihood estimation problems. Instead of completely maximizing the likelihood function in each of the iterations of the GEM algorithm, the algorithm chooses a set of parameters that ensures the likelihood's value increases with each iteration. As such, the GEM algorithm satisfies the key condition for convergence to the optimum presented in Casella and Berger (2001), Theorem 7.2.20. However, it is possible for the GEM algorithm to fail to converge, as I do not establish formal almost sure convergence results for the algorithm as presented for a class of general stochastic optimizaiton procedures in Biscarat (1994), Chan and Ledolter (1995) and Sherman, Dalal, and Ho (1999).

Appendix 1.A.5: Level-k Bidding with Reserve Prices

This appendix develops two key result regarding level-k bidding with reserve prices that closely parallel existing results for equilibrium behavior. First, it presents the general level-k bidding strategy when there is a reserve price in the auction. Second, it characterizes the effect of uncertain competition on level-k bidding behavior. The appendix closes with a result characterizing the expected revenue in a level-k auction as a weighted average of the expected revenues conditional on the composition of the bidding population.

1.A.5.1 Certain Competition

When bidders know the number of participating bidders in the auction, counterfactual bidding behavior treats the distribution over valuations for participating bidders conditions on the valuation being greater than the reserve price. As such, denoting the level $(k-1)$ bidding strategy when the reserve price is r

by $\sigma_{L(k-1)_\tau,r}(X)$ the inverse bidding function from equation 1.12 in section 1.4.2 incorporates this information:

$$X = \sigma_{Lk_\tau}(X) + \frac{F_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right) - F_X(r)}{f_X\left(\sigma^{-1}_{L(k-1)_\tau}\left(\sigma_{Lk_\tau}(X)\right)\right)\frac{d\sigma^{-1}_{L(k-1)_\tau}(s)}{ds}|_{s=\sigma_{Lk_\tau}(X)}} \tag{1.A.13}$$

As in equilibrium, the effect of the reserve price on level-k bidding behavior is to shift the differential equation defining bid shades to initialize at the reserve price rather than the minimum valuation (which is here set to zero for exposition). This effect is apparent in the estimated bidding strategies displayed in figure 8's depiction based on estimates from the USFS timber auction data.

1.A.5.2 Uncertain Competition

When bidders know only the number of potentially participating bidders, $\overline{N}$, but not the actual number of potential bidders with valuations that exceed the reserve price, N, they face an uncertain amount of competition in the auction. In another parallel to a well-known equilibrium result stated in Krishna (2002), the bidding strategy for level-k bidders with uncertain competition is a weighted average of the bidding strategies with a fixed number of participating bidders. To establish this result, first suppose the number of bidders varies exogenously and denote $Pr\{N = n\} = q_n$. Then given the level $(k-1)$ bidding strategy, the expected utility from the level-k bidder-type's bid is:

$$E\left[U\left(X_i, s_1, \ldots, s_N\right)|X_i\right] = \sum_{n=1}^{\overline{N}} \left(X_i - s_i\right) q_n Pr\left\{s_i > s_{-i}\right\}^{n-1}$$

Substituting $Pr\{s_i > s_{-i}\} = F_X\left(\sigma^{-1}_{L(k-1)_\tau}(s_i)\right)$ and taking first order conditions gives the inverse bidding function as:

$$\begin{aligned} X &= \sum_{n=1}^{\overline{N}} \frac{(n-1)\, q_n F_X\left(\sigma_{L(k-1)_\tau}^{-1}(s_i)\right)}{\sum_{m=1}^{\overline{N}} (m-1)\, q_m F_X\left(\sigma_{L(k-1)_\tau}^{-1}(s_i)\right)} \cdot \\ &\quad \left\{ s_i + \frac{F_X\left(\sigma_{L(k-1)_\tau}^{-1}(s_i)\right)}{(n-1)\, f_X\left(\sigma_{L(k-1)_\tau}^{-1}(s_i)\right) \frac{d\sigma_{L(k-1)_\tau}^{-1}(\xi)}{d\xi}\Big|_{\xi=s_i}} \right\} \\ &= \sum_{n=1}^{\overline{N}} \omega_n(s_i)\, \sigma_{k_n}^{-1}(s_i) \end{aligned} \tag{1.A.14}$$

In the case of a binding reserve price, q_n is the probability that n bidders will have a valuation exceeding the reserve price given that player i's valuation is above the reserve price, which is given by:

$$q_n = \binom{\overline{N}-1}{n-1} F_X(r)^{N-n} \left(1 - F_X(r)\right)^{n-1}$$

So the inverse bidding function with endogenous participation is:

$$X = \sum_{n=1}^{\overline{N}} \omega_{n,r}(s_i)\, \sigma_{k_n,r}^{-1}(s_i)$$

1.A.5.3 Expected Revenues

First, consider the seller's expected revenue when participating bidders know the number of bidders participating in each auction. In this case, the cdf for the winning bid is given by F_{W_N} in equation 1.17. Ex ante, the seller does not know the number of bidders whose valuations will exceed the reserve price but only the number of potential bidders, $\overline{N}$. Accounting for this uncertainty, the cdf for the seller's revenue is then:

$$F_{W_{\overline{N}}}(w;\theta) = \sum_{N=0}^{\overline{N}} \binom{\overline{N}}{N} F_X(r)^{\overline{N}-N} \left(1 - F_X(r)\right)^N F_{W_N}(w;\theta) \tag{1.A.15}$$

The key feature of this formula for analyzing expected revenues is that, since F_{W_N} is a continuous polynomial in the distribution over types, the expected revenue to the seller is continuous in the distribution over types. Further, because the behavior of each bidder and bidder-type is independent of one another, the distribution over bidder-types that maximizes and minimizes these expected revenues corresponds to the homogeneous populations of bidder-types that individually maximize and minimze the sellers' expected revenue, respectively. This result is stated in the following theorem:

Theorem 1.9 (Identified Set for Expected Revenues)
Suppose the distribution over bidder-types is not identified in the level-k auction model. The seller's expected revenue at a given reserve price belongs to a closed, convex identified set. Further, this identified set is bounded above and below by the expected revenues generated by the homogeneous population of bidder-types that maximize and minimize expected revenues, respectively.

Bibliography

Victor Aguirreagabiria and Arvind Magesan. Estimation of dynamic discrete games when players beliefs are not in equilibrium. Mimeo, 2009.

Yonghong An, Yingyao Hu, and Matthew Shum. Estimating first-price auction models with unknown number of bidders: a misclassification approach. Mimeo, 2009.

Andres Aradillas-Lopez and Elie Tamer. The identification power of equilibrium in simple games. Journal of Business & Economic Statistics, 26(3):261–283, JUL 2008. ISSN 0735-0015.

Zvi Artstein and Roger J.-B. Wets. Consistency of minimizers and the slln for stochastic programs. Journal of Convex Analysis, 2(1-2):1–17, 1995.

S Athey and PA Haile. Identification of standard auction models. Econometrica, 70(6):2107–2140, NOV 2002. ISSN 0012-9682.

S Athey and J Levin. Information and competition in us forest service timber auctions. Journal of Poltiical Economy, 109, 2001.

Susan Athey and Philip A. Haile. Handbook of Econometrics, chapter Nonparametric Approaches to Auctions. Elsevier, 2005.

Susan Athey, Jonathan Levin, and Enrique Seira. Comparing open and sealed bid auctions: Evidence from timber auctions. NBER Working Papers, 14590, 2008.

P Bajari and A Hortacsu. Are structural estimates of auction models reasonable? Evidence from experimental data. Journal of Political Economy, 113(4):703–741, AUG 2005. ISSN 0022-3808.

P Bajari and LX Ye. Deciding between competition and collusion. Review of Economics and Statistics, 85(4):971–989, NOV 2003. ISSN 0034-6535.

Patrick Bajari, Victor Chernozhukov, Han Hong, and Denis Nekipelov. Semiparametric estimation of a dynamic game of incomplete information. Mimeo, 2008.

LH Baldwin, RC Marshall, and JF Richard. Bidder collusion at forest service timber sales. Journal of Political Economy, 105, 1997.

P Battigalli and M Siniscalchi. Rationalizable bidding in first-price auctions. Games and Economic Behavior, 45(1):38–72, OCT 2003. ISSN 0899-8256.

C. Lanier Benkard and Steven Berry. On the nonparametric identification of non-linear simultaneous equations models: Comment on Brown (1983) and Roehrig (1988). Econometrica, 74(5):1429–1440, SEP 2006. ISSN 0012-9682.

Arie Beresteanu and Francesca Molinari. Asymptotic properties for a class of partially identified models. Econometrica, 76(4):763–814, JUL 2008. ISSN 0012-9682.

BD Bernheim. Rationalizable Strategic Behavior. Econometrica, 52(4):1007–1028, 1984. ISSN 0012-9682.

ST Berry. Estimation of a Model of Entry in the Airline Industry. Econometrica, 60(4):889–917, JUL 1992. ISSN 0012-9682.

Herman J. Bierens. Semi-nonparametric interval-censored mixed proportional hazard models: Identification and consistency results. Mimeo, (2003-09), September 2006.

Herman J. Bierens and Hosin Song. Semi-nonparametric estimation of independently and identically repeated first-price auctions via simulated method of moments. Mimeo, (2003-09), September 2007.

JC Biscarat. Almost Sure Convergence of a Class of Stochastic Algorithms. Stochastic Processes and Their Applications, 50(1):83–99, MAR 1994. ISSN 0304-4149.

A Bosch-Domenech, JG Montalvo, R Nagel, and A Satorra. One, two, (three), infinity, ... : Newspaper and lab beauty-contest experiments. American Economic Review, 92(5):1687–1701, DEC 2002. ISSN 0002-8282.

Bjarne Brendstrup and Harry J. Paarsch. Identification and estimation in sequential, asymmetric, english auctions. Journal of Econometrics, 134(1):69–94, September 2006.

TF Bresnahan and PC Reiss. Entry and Competition in Concentrated Markets. Journal of Political Economy, 99(5):977–1009, OCT 1991. ISSN 0022-3808.

BW Brown. The Identification Problem in Systems Non-linear in the Variables. Econometrica, 51(1):175–196, 1983. ISSN 0012-9682.

CF Camerer, TH Ho, and JK Chong. A cognitive hierarchy model of games. Quarterly Journal of Economics, 119(3):861–898, AUG 2004. ISSN 0033-5533.

S Campo, E Guerre, I Perrigne, and Q Vuong. Semiparametric estimation of first-price auctions with risk averse bidders. Mimeo, 2000.

Joeorge Casella and Roger Berger. Statistical Inference. Duxbury Press, 2001.

KS Chan and J Ledolter. Monte-Carlo EM Estimation for Time-Series Models Involving Counts. Journal of the American Statistical Association, 90(429):242–252, MAR 1995. ISSN 0162-1459.

Xiaohong Chen. Large sample sieve estimation of semi-nonparametric models. Handbook of Econometrics, 6B:5549–5632, 2007.

V Chernozhukov and A Hong. Likelihood estimation and inference in a class of nonregular econometric models. Econometrica, 72(5):1445–1480, SEP 2004. ISSN 0012-9682.

Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. Econometrica, 75(5):1243–1284, SEP 2007. ISSN 0012-9682.

Federico Ciliberto and Elie T. Tamer. Market structure and multiple equilibria in the airline industry. Econometrica, Forthcoming 2009.

Allan Collard-Wexler. Discussion of aradillas-lopez and tamer. Journal of Business and Economic Statistics, 2008.

M Costa-Gomes, VP Crawford, and B Broseta. Cognition and behavior in normal-form games: An experimental study. Econometrica, 69(5):1193–1235, SEP 2001. ISSN 0012-9682.

Miguel A. Costa-Gomes and Vincent P. Crawford. Cognition and behavior in two-person guessing games: An experimental study. American Economic Review, 96(5):1737–1768, DEC 2006. ISSN 0002-8282.

JC Cox, VL Smith, and JM Walker. Theory and Misbehavior of 1st-Price Auctions - Comment. American Economic Review, 82(5):1392–1412, DEC 1992. ISSN 0002-8282.

V. Crawford, T. Kugler, Z. Neeman, and A. Pauzner. Behaviorally optimal auction design: An example and some observations. Journal of the European Economic Association, 7, 2009.

Vincent Crawford and Nagore Iriberri. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? Econometrica, 75(6):1721–1770, NOV 2007a. ISSN 0012-9682.

Vincent P. Crawford and Nagore Iriberri. Fatal attraction: Salience, naivete, and sophistication in experimental "Hide-and-Seek" games. American Economic Review, 97(5):1731–1750, DEC 2007b. ISSN 0002-8282.

VP Crawford. Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. American Economic Review, 93(1):133–149, MAR 2003. ISSN 0002-8282.

SG Donald and HJ Paarsch. Identification, estimation, and testing in parametric empirical models of auctions within the independent private values paradigm. Econometric Theory, 12(3):517–567, AUG 1996. ISSN 0266-4666.

MA El Gamal and DM Grether. Are people Bayesian? Uncovering behavioral strategies. Journal of the American Statistical Association, 90(432):1137–1145, DEC 1995. ISSN 0162-1459.

D Ellsberg. Risk, Ambiguity, and the Savage Axioms. Quarterly Journal of Economics, 75(4):643–669, 1961. ISSN 0033-5533.

E Eyster and M Rabin. Cursed equilibrium. Econometrica, 73(5):1623–1672, SEP 2005. ISSN 0012-9682.

Philippe Fevrier. Nonparametric identification and estimation of a class of common value auction models. Journal of Applied Econometrics, 23(7):949–964, NOV-DEC 2008. ISSN 0883-7252.

I Gilboa and D Schmeidler. Max-min expected utility with non-unique prior. Journal of Mathematical Economics, 18, 1989.

JK Goeree, CA Holt, and TR Palfrey. Quantal response equilibrium and overbidding in private-value auctions. Journal of Economic Theory, 104(1):247–272, MAY 2002. ISSN 0022-0531.

JK Goeree, CA Holt, and TR Palfrey. Regular Quantal Response Equilibrium. Experimental Economics, 8(4):347–367, DEC 2005. ISSN 1386-4157.

E Guerre, I Perrigne, and Q Vuong. Optimal nonparametric estimation of first-price auctions. Econometrica, 68(3):525–574, MAY 2000. ISSN 0012-9682.

Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Nonparametric Identification of Risk Aversion in First-Price Auctions Under Exclusion Restrictions. Econometrica, 77(4):1193–1227, JUL 2009. ISSN 0012-9682.

P Haile. Auctions with resale markets: An application to us forest service timber sales. American Economic Review, 91, 2001.

PA Haile and E Tamer. Inference with an incomplete model of English auctions. Journal of Political Economy, 111(1):1–51, FEB 2003. ISSN 0022-3808.

Philip A. Haile, Han Hong, and Matthew Shum. Nonparametric tests for common values at first-price sealed-bid auctions. SSRN eLibrary, 2003.

Philip A. Haile, Ali Hortacsu, and Grigory Kosenok. On the empirical content of quantal response equilibrium. American Economic Review, 98(1):180–200, MAR 2008. ISSN 0002-8282.

C Han and PCB Phillips. GMM with many moment conditions. Econometrica, 74 (1):147–192, JAN 2006. ISSN 0012-9682.

Daniel Henderson, John List, Daniel Millimet, Christopher Parmeter, and Michael Price. Imposing monotonicity nonparametrically in first-price auctions. Mimeo, 2008.

C Hess. Epi-convergence of sequences of normal integrands and strong consistency of the maximum likelihood estimator. Annals of Statistics, 24(3):1298–1315, JUN 1996. ISSN 0090-5364.

TH Ho, C Camerer, and K Weigelt. Iterated dominance and iterated best response in experimental "p-beauty contests". American Economic Review, 88(4):947–969, SEP 1998. ISSN 0002-8282.

H Hong and M Shum. Increasing competition and the winner's curse: Evidence from procurement. Review of Economic Studies, 69(4):871–898, OCT 2002. ISSN 0034-6527.

Yingyao Hu, David McAdams, and Matthew Shum. Nonparametric identification of auction models with non-separable unobserved heterogeneity. Mimeo, 2009.

Asen Ivanov, Dan Levin, and Muriel Niederle. Can relaxation of beliefs rationalize the winners curse?: An experimental study. Mimeo, 2008.

J Kagel and A Roth. Handbook of Experimental Economics. Princeton University Press, 1997.

Hiroaki Kaido. A dual approach to inference for partially identified econometric models. Mimeo, 2009.

P Klibanoff, M Marinacci, and S Mukerji. A smooth model of decision making under ambiguity. Econometrica , 73(6):1849–1892, NOV 2005. ISSN 0012-9682.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive smooth ambiguity preferences. Journal of Economic Theory, 144(3):930–976, MAY 2009. ISSN 0022-0531.

Elena Krasnokutskaya. Identification and estimation in procurement auctions under unobserved auction heterogeneity. Review of Economic Studies, Forthcoming - 2009.

Vijay Krishna. Auction Theory. Academic Press, 2002.

JJ Laffont and Q Vuong. Structural econometric analysis of descending auctions. European Economic Review, 37, 1993.

JJ Laffont, H Ossard, and Q Vuong. Econometrics of first price auctions. Econometrica, 63, 1995.

Steven D. Levitt and John A. List. What do laboratory experiments measuring social preferences reveal about the real world? Journal of Economic Perspectives, 21(2):153–174, SPR 2007. ISSN 0895-3309.

T Li. Econometrics of first-price auctions with entry and binding reservation prices. Journal of Econometrics, 126, 2005.

T Li and I Perrigne. Timber sale auctions with random reserve prices. Review of Economics and Statistics, 85(1):189–200, FEB 2003. ISSN 0034-6535.

John A. List and David Lucking-Reiley. Demand reduction in a multi-unit auction: Evidence from a sportscard field experiment. American Economic Review, 90 (4):961–972, September 2000.

J Lu. Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. Mimeo, 2004.

Jingfeng Lu and Isabelle Perrigne. Estimating Risk Aversion from Ascending and Sealed-Bid Auctions: The Case of Timber Auction Data. Journal of Applied Econometrics, 23(7):871–896, NOV-DEC 2008. ISSN 0883-7252.

Xun Lu. Testing for treatment dependence and heterogeneity of effects in nonseparable structural equations. Mimeo, 2009.

D Lucking-Reiley. Using field experiments to test equivalence between auction formats: Magic on the Internet. American Economic Review, 89(5):1063–1080, DEC 1999. ISSN 0002-8282.

Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. Econometrica, 74 (6):1447–1498, NOV 2006. ISSN 0012-9682.

Charles Manski. Identification Problems in the Social Sciences. Harvard University Press, 1995.

E Maskin and J Riley. Optimal auctions with risk averse bidders. Econometrica, 52 (6), 1984.

Rosa L. Matzkin. Identification in nonparametric simultaneous equations models. Econometrica, 76(5):945–978, SEP 2008. ISSN 0012-9682.

RD McKelvey and TR Palfrey. Quantal Response Equilibria for Normal-Form Games. Games and Economic Behavior, 10(1):6–38, JUL 1995. ISSN 0899-8256.

RD McKelvey and TR Palfrey. Quantal Response Equilibria for Extensive Form Games. Experimental Economics, 10(1):6–38, JUL 1998. ISSN 0899-8256.

PR Milgrom and RJ Weber. A Theory of Auctions and Competitive Bidding. Econometrica, 50(5):1089–1122, 1982. ISSN 0012-9682.

RB Myerson. Optimal Auction Design. Mathematics of Operations Research, 6 (1):58–73, 1981. ISSN 0364-765X.

R Nagel. Unraveling in guessing games: An experimental study. American Economic Review, 85(5):1313–1326, DEC 1995. ISSN 0002-8282.

R Paarsch and H Hong. An Introduction to the Structural Econometrics of Auction Data. The MIT Press, 2006.

DG Pearce. Rationalizable Strategic Behavior and the Problem of Perfection. Econometrica, 52(4):1029–1050, 1984. ISSN 0012-9682.

Isabelle Perrigne and Quang Vuong. Identification and estimation of Bidders' risk aversion in first-price auctions. American Economic Review, 97(2):444–448, MAY 2007. ISSN 0002-8282.

A Rapoport, DA Seale, and E Winter. Coordination and learning behavior in large groups with asymmetric players. Games and Economic Behavior, 39(1): 111–136, APR 2002. ISSN 0899-8256.

CS Roehrig. Conditions for identificaiton in nonparametric and parametric models. Econometrica, 56(2):433–447, MAR 1988. ISSN 0012-9682.

Brian W. Rogers, Thomas R. Palfrey, and Colin F. Camerer. Heterogeneous quantal response equilibrium and cognitive hierarchies. Journal of Economic Theory, 144(4):1440–1467, JUL 2009. ISSN 0022-0531.

Leonard J Savage. The Foundations of Statistics. Dover Publications, 1954.

Robert Sherman, Sid Dalal, and Yu-Yun Ho. Conditions for convergence of montecarlo em sequences with an application to diffusion modeling. Econometrics Journal, 2:248–267, 1999.

DO Stahl and PW Wilson. Experimental Evidence on Players Models of Other Players. Journal of Economic Behavior & Organization, 25(3):309–327, DEC 1994. ISSN 0167-2681.

DO Stahl and PW Wilson. On Players Models of Other Players - Theory and Experimental Evidence. Games and Economic Behavior, 10(1):218–254, JUL 1995. ISSN 0899-8256.

Chapter 2

Bayesian Methods for Covariance Matrix Estimation

This paper proposes Bayesian methods adapted to estimating a covariance matrix for a large number of random variables. The analysis models the return generating process with an unrestricted factor model of covariance, imposing structure on the covariance matrix through prior beliefs on the parameters governing this unrestricted return generating process. By nesting many popular shrinkage estimators for covariance matrices, these results provide an Empirical Bayesian foundation for a general class of shrinkage estimators and use the shrinkage interpretation to characterize prior beliefs that optimize a posterior objective function. The consistent estimator coupled with economically motivated priors delivers lower finite-sample loss than existing estimators in Monte Carlo simulations and performs well in applied settings, as illustrated in a minimum variance portfolio selection exercise.

2.1 Introduction

A variety of financial and economic problems focusing on variance reduction require analyzing a large covariance matrix that may be difficult to estimate with precision. The curse of dimensionality presents the key challenge to this estimation problem as the covariance matrix for N random variables representing asset returns has $\frac{N(N-1)}{2}$ free parameters that must be estimated with only T observations where T is often less than N. The unbiased sample covariance matrix estimator in this setting is extremely noisy and necessarily has $N - T$ zero eigenvalues and a zero determinant, rendering the estimated covariance matrix non-invertible and, consequently, yielding non-unique and highly unstable solutions to variance minimization problems.

Traditional approaches to this problem assume a risk-factor model to the covariances between assets to impose structure for the covariance matrix by isolating the sources of systematic risk without restricting the variances of each random variable. While these models greatly improve the precision of the estimated covariance matrix, this increased precision comes at the cost of a misspecified model and an inconsistent estimator. Further, this approach requires a model selection exercise to identify the risk factors that characterize systematic risk and leaves still the problem of estimating the covariance matrix for these systematic risk factors.

In order to smooth the trade-off between bias and variance in traditional models, pseudo-Bayesian shrinkage methods treat the factor-based model as if it were an investor's prior belief for the covariance matrix. The main tool in this nascent literature is presented in a series of papers by Ledoit and Wolf (2004a,b), who propose a James-Stein shrinkage estimator corresponding to a weighted average of the sample covariance matrix and the risk-factor model of covariance. In their analysis, Ledoit & Wolf characterize the optimal weights that minimize the finite-sample mean squared error of the estimator. In this way, the Ledoit & Wolf shrinkage estimator is, by definition, admissible (under the specified loss function) in the set of shrinkage estimators for any fixed set of risk factors. While their estimator works well in application, Ledoit & Wolf don't present a formal analysis of the distributions relating to prior beliefs and the likelihood of the data to give their shrinkage estimator a formal grounding in Bayesian estimation.

In this paper, I present a truly Bayesian approach to this problem that incorporates the investor's belief in a risk-factor model as a prior distribution over factor loadings in a Bayesian regression framework. This approach is similar to Bayesian methods for incorporating theoretical pricing models as a prior belief when estimating expected returns, characterizing the prior belief as a restricted model nested within a larger regression framework. The key methodological innovation augments the K risk factors with a set of $(N-K)$ derived factors that allow for unmodelled sources of covariance across the random variables. In the augmented regression context, traditional risk-factor models of covariance can be thought of as the posterior covariance matrix in a Bayesian model where the prior specifies that factor loadings on the derived factors augmenting the original K factors are believed to be zero with certainty. In another departure from most of the Bayesian asset pricing literature, these informative priors are not solely related to the "alpha" in the return generating process, but also the "betas" themselves. In this sense, the model incorporates beliefs relating to the covariances across returns.

The Bayesian analysis yields an intuitive estimator for covariance matrices that can be readily coupled with Bayesian methods for estimating means, providing a unified Bayesian framework for jointly analyzing expectations, variances, and covariances. Under a conjugate prior specification, I present a closed-form solution for the posterior factor loadings that is the usual matrix-weighted average of prior expected factor loadings and the ordinarly least squares-estimated factor loadings. The closed form solution avoids challenges in sampling from the extremely high-dimensional space of covariances under more general distributional assumptions, presenting an easily computable posterior covariance matrix with attractive analytical properties. Further separability conditions on the factors and prior beliefs admit a represention of the posterior covariance matrix as a multi-factor shrinkage estimator, providing an empirical Bayesian foundation for a very broad class of shrinkage estimators for covariance matrices, including Ledoit & Wolf's approach.

Bayesian methods for estimating covariance matrices provide a flexible means of introducing structure to covariance matrix estimation that admits economically motivated priors based on empirical regularities in analyzing equity data. I evaluate the empirical value of these methods using three separate prior specifications. The first specification of prior beliefs naturally represents the belief in a single factor model of covariance, introducing informative beliefs that factor loadings on any other factors are zero. The second specification of prior beliefs represents the empirical regularity that estimated factor loadings tend to be mean-reverting. In this specification, prior beliefs for the factor loadings of a given security are centered at the grand average of factor loadings taken across all securities. Lastly, I follow Ledoit & Wolf's analysis to introduce a set of "optimal" prior beliefs that minimize a finite sample expected loss for the posterior estimated covariance matrix.

The empirical evidence suggests that these Bayesian techniques are extremely useful directly in estimating covariance matrices as well as in applied portfolio selection exercises. In Monte Carlo simulations, the Bayesian covariance matrices deliver lower finite sample expected loss than existing estimation techniques, especially in settings where the number of observations is small relative to the dimension of the covariance matrix. Similarly, when applied to estimating minimum variance portfolios, backtest and simulation tests illustrate that the Bayesian estimation strategy delivers portfolios that generate lower variance than existing techniques. Beyond providing the theoretical basis for shrinkage estimators, a main contribution of this work is to introduce estimators providing even better performance than these techniques.

The paper proceeds by discussing the analytical properties of the posterior covariance matrix in section 2.2, which formally describes the statistical model, presents the closed form posterior covariance matrix in the natural-conjugate setting, and establishes some analytical properties of the posterior covariance matrix, including consistency for the true covariance matrix and a shrinkage decomposition of the posterior covariance matrix under an orthogonal factor structure. Section 2.3 characterizes the finite-sample loss measures and solves for optimal priors before section 2.4 reviews more economically motivated candidate priors. I analyze the finite-sample performance of the estimator using Monte Carlo simulation evidence in section 2.5, illustrating that the bayesian approach is empirically valuable and presents an improvement over existing estimators in many settings. Section 2.6 illustrates the utility of the posterior covariance matrix in portfolio selection, where the minimum variance porfolios estimated using the Bayesian Factor Covariance estimator provide dramatic improvements over alterative sample and shrinkage-based estimators. Section 2.7 concludes by discussing additional applications and potential extensions to the current work.

2.2 Bayesian Estimator for Covariance Matrices

This section develops the statistical model and derive the posterior expectations for covariance matrices in a natural conjugate setting. The key to the Bayesian analysis of the covariance matrix lies in representing the unrestricted covariance matrix as an unrestricted N-factor model of covariance. This representation allows for structure in the posterior covariance matrix by imposing the restrictions from a structured factor model through prior beliefs. Further, the structure of the posterior covariance matrix provides the key to characterizing prior beliefs consistent with existing shrinkage estimators.

2.2.1 Statistical Model

The objective is to estimate the covariance matrix for the returns on N securities, $r_1, \ldots, r_N$, each of which are normally distributed with known means $\mu = [\mu_1, \ldots, \mu_N]'$ and an unknown covariance matrix Σ. Assume that there are K benchmarks and economic factors $F_1, \ldots, F_K$ that represent systematic sources of variance across the securities, and that these factors have known covariance matrix Γ_F. Using principal components analysis, the residuals from the regression of security returns on the systematic factors can be transformed into $N - K$ augmenting factors, $F_{K+1}, \ldots, F_N$ with diagonal covariance matrix Γ_R. The present analysis ignores any estimation error in deriving these factors or recovering their covariance matrix Γ_R.

Denoting the full set of N factors driving covariance by $F = [F_1, \ldots, F_N]$, represent the return generating process for asset i in period t without loss of generality as:

$$r_{i,t} = \alpha_i + \sum_{k=1}^{N} \beta_{i,k} F_{i,t} + \epsilon_{i,t} \tag{2.1}$$
$$= \alpha_i + \beta_i' F_t + \epsilon_{i,t}$$

In this return generating process, the vector $\beta_i = [\beta_{i,1}, \ldots, \beta_{i,N}]'$ represents the factor loadings for asset i and the error term ϵ_i represents idiosyncratic noise in asset i's return that is uncorrelated with any of the N factors and uncorrelated with other assets' idiosyncratic noise. Since the returns for asset i are fully explained by the set of factors, there is no idiosyncratic variation and $\epsilon_i = 0$. However, this analysis treats ϵ_i as having a non-degenerate normal distribution with variance $\sigma^2_{\epsilon,i}$, as though white noise were added to asset i's return series after the factors have been extracted. As is described in greater detail below, the magnitude of this variance can then be interpreted as a bandwidth parameter for the estimator.

The unrestricted covariance matrix implied by equation 2.1's return generating process takes the usual diagonalizable form with B denoting a matrix containing the factor loadings for all securities, Γ denoting the covariance matrix for the factors, and Λ denoting the diagonal matrix of idiosyncratic variances:

$$\Sigma = B\Gamma B' + \Lambda, \tag{2.2}$$

$$\text{where, } B = \begin{bmatrix} \beta_1' \\ \beta_2' \\ \vdots \\ \beta_N' \end{bmatrix}, \Gamma = \begin{bmatrix} \Gamma_F & 0 \\ 0 & \Gamma_R \end{bmatrix}, \text{ and, } \Lambda = \begin{bmatrix} \sigma^2_{\epsilon,1} & 0 & \cdots & 0 \\ 0 & \sigma^2_{\epsilon,2} & \cdots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \sigma^2_{\epsilon,N} \end{bmatrix}.$$

Factor models impose structure on the covariance matrix by implicitly restricting a subset of the factor loadings (typically those associated with the $N-K$ derived factors) in the return generating process from equation 2.1 to equal zero. The alternative to this threshold-type restriction frames the factor model as the prior belief within a Bayesian regression framework. The next section will discuss specific priors in greater detail, but for now it suffices to represent the investor's prior beliefs about as:

$$\beta_i, \sigma^2_{\epsilon,i} \sim_{\text{prior}} NG\left(\beta_{i,0}, \Omega_{i,0}, v_{i,0}, s^2_{i,0}\right) \tag{2.3}$$

Here "NG" represents the normal-gamma distribution so β_i has a Normal prior distribution with mean $\beta_{i,0}$ and covariance matrix Ω and the idiosyncratic variance $\sigma^2_{\epsilon,i}$ has an independent Gamma distribution with $v_{i,0}$ degrees of freedom and location parameter $s^2_{i,0}$.

Given T observations from the normal return generating process in equation 2.1, the likelihood of the data for specific values of β_i and $\sigma^2_{\epsilon,i}$ is given by a conditional Normal-Gamma distribution. That is, the likelihood for the true β_i corresponds to a normal distribution with expectation given by the OLS estimates of factor loadings, $\hat{\beta}$, and covariance matrix $\sigma^2_{\epsilon,i}F'F$ conditional on $\sigma^2_{\epsilon,i}$, which has an unconditional gamma distribution with $T - N$ degrees of freedom and location parameter s^2, the OLS-computed standard error of residuals.

$$\tilde{p}\left(R|\beta, \sigma^2_{\epsilon,i}\right) = N\left(\hat{\beta}, \sigma^2_{\epsilon,i}F'F\right), \text{ and, } \tilde{p}\left(\sigma^2_{\epsilon,i}\right) = G\left(T - N, s^2\right) \tag{2.4}$$

Now, since $s^2 = 0$ in the sample, the likelihood above is not well-defined. This singularity occurs because the data is perfectly described by the model, an event that also arises in non-parametric regression. To address this overfitting, introduce additional noise to each security's return that prevents the factors from perfectly explaining each asset's return. The variance of this noise, $\frac{h^2}{T}$ can be interpreted as the bandwidth of the covariance matrix estimator, which is scaled by the sample size to ensure estimator consistency. The likelihood for the Bayesian analysis is then:

$$p\left(R|\beta, \sigma^2_{\epsilon,i}\right) = N\left(\hat{\beta}, \sigma^2_{\epsilon,i}F'F\right), \text{ and, } p\left(\sigma^2_{\epsilon,i}\right) = G\left(T - N, s^2 + \frac{h^2}{T}\right) \tag{2.5}$$

With this likelihood, the prior specification and likelihood correspond to a natural conjugate setting, yielding analytical posterior expectations for each asset's factor loadings in closed-form. As is established in textbook treatments on

Bayesian econometrics such as Koop (2003) or Geweke (2005), the posterior expected factor loadings are the matrix-weighted average of the prior factor loadings and the OLS estimated factor loadings:

$$\beta_i^* \equiv E_{post}\left[\beta_i\right] = \left(\Sigma_0^{-1} + F'F\right)^{-1}\left(\Sigma_0^{-1}\beta_{i,0} + F'F\hat{\beta}_i\right) \tag{2.6}$$

Also, the posterior expected idiosyncratic variance ($E_{post}\left[\sigma_{\epsilon,i}^2\right]$, which is denoted s_i^{*2}) is given by a weighted average of the prior expected idiosyncratic variance, the sample idiosyncratic variance, and a term that captures the disparity between the prior and OLS factor loadings:

$$\begin{aligned}(T + v_{i,0})\, s_i^{*2} =& v_{i,0}s_{i,0}^2 + (T - N)\left(s_i^2 + \frac{h^2}{T}\right) \\ &+ \left(\hat{\beta}_i - \beta_i^*\right)' F'F\left(\hat{\beta}_i - \beta_i^*\right) + \left(\beta_{i,0} - \beta_i^*\right)' \Omega^{-1}\left(\beta_{i,0} - \beta_i^*\right)\end{aligned} \tag{2.7}$$

Defining the matrices B^* and Λ^* as the posterior expectations for the matrices B and Λ defined above, the posterior expectation for the covariance matrix is:

$$\Sigma^* = B^*\Sigma_F B^{*'} + \Lambda \tag{2.8}$$

As is common with Bayesian estimators, as the amount of information in the data dwarfs the prior belief, the posterior expectation converges to the unbiased sample estimator. This convergence ensures that the estimator will be asymptotically consistent for the true covariance matrix.

Theorem 2.1 *The posterior covariance matrix estimator is consistent:*

$$\mathrm{p}\lim_{x\to\infty} \Sigma^* = \Sigma \tag{2.9}$$

Proof. From equation (4), it's clear that $\mathrm{plim}_{x\to\infty}\, \beta^* = \mathrm{plim}_{x\to\infty}\, \hat{\beta}_i = \beta_i$. This convergence implies that $\mathrm{plim}_{x\to\infty}\, B^* = \mathrm{plim}_{x\to\infty}\, \hat{B} = B$ and so, since Σ_F and Λ are known (the latter, given B), the result holds. ■

2.2.2 Empirical Bayesian Foundations for Shrinkage Estimators

To further characterize the properties of the posterior covariance matrix, consider the special setting when factors and beliefs are orthogonal. Here $\Omega_{i,0} = \sigma_C^2 I_N$ and Σ_F is a diagonal matrix with the i-th entry $\sigma_{F_i}^2$. In this case, the posterior factor loadings for each factor correspond to a weighted average between the prior expected factor loading and the estimated factor loading where the weight assigned to the prior expectation is constant across all assets. This feature of the estimator admits a shrinkage interpretation by writing the posterior covariance matrix as a weighted sum of $2N$ single-factor models, providing a common frame with existing shrinkage estimators. To establish the shrinkage decomposition, denote by $\delta_k = \frac{\sigma_C^2}{\sigma_C^2 + T\sigma_{F_k}^2}$ the weight assigned to the estimated factor k loading, $B_{0,k}$ the $N \times 1$ vector of each asset's prior expected k factor loadings, and $\hat{B}_k$ the vector of each asset's estimated k factor loadings. Then the posterior covariance matrix can be written as:

$$\Sigma^* = B^* \Sigma_F B^{*'} + \Lambda = \sum_{k=1}^{N} \delta_k \sigma_{F_k}^2 \hat{B}_k \hat{B}_k' + \sum_{k=1}^{N} (1 - \delta_k)\, \sigma_{F_k}^2 B_{0,k} B_{0,k}' + \Lambda \qquad (2.10)$$

This setting allows a visual presentation of the convergence of the posterior covariance matrix to the sample covariance matrix by analyzing the weight the posterior assigns to the prior factor loading for each of the factors based on the updating formula in equation 2.6. Figure 2.1 illustrates the convergence as data accumulates for a five asset covariance matrix where all factors are derived using principal components and the prior is specified with $\beta_{i,0,k} = 0$, $\Omega_{i,0} = I_N$, $v_{i,0} = 0$, and $s_{i,0}^2 = 1, \forall i, k$. As is immediately apparent, the first factor, with the highest variance, shrinks towards the sample factor loading quite quickly as the number of observations increases and the other factors follow in succession. Still, it is interesting to note that, even in a sample with over 10,000 observations, the prior factor loading for the last factor still receives over 25% weight in the posterior weighting, indicating the importance of the prior for reining in factor loadings that would otherwise be very imprecisely estimated.

Another benefit of the decomposition in 2.10 is that it facilitates characterizing prior beliefs consistent with a given posterior expectation. The next section turns to the analytical exercise of finding prior beliefs consistent with existing shrinkage estimators, beginning with the Ledoit and Wolf (2004a) estimator. This analysis illustrates a general algorithm for deriving prior beliefs consistent with shrinkage of the sample covariance matrix towards any positive-semidefinite prior covariance matrix, allowing an empirical Bayesian foundation for the shrinkage technique proposed by Jagannathan and Ma (2003), who represent non-negativity constraints as arising from a shrinkage of the variance-covariance matrix.

Empirical Bayesian Priors for Ledoit & Wolf Shrinkage

The main result in this section provides an Empirical Bayesian procedure that leads to the Ledoit & Wolf Single-Factor Shrinkage estimator where the prior variance around factor loadings will depend on the variance of the factor itself. In addition to providing a truly Bayesian interpretation of the Ledoit & Wolf shrinkage estimator, the analysis illustrates the appeal of their model as a particularly parsimonious specification for prior beliefs. As background, denote the Ledoit & Wolf shrinkage estimator as:

$$\begin{aligned} \Sigma^*_{LW} &= (1-\delta)\,\Sigma_{SF} = \delta\Sigma_S \\ &= (1-\delta)\left(B_{SF}\sigma^2_{SF}B'_{SF} + \Lambda_{SF}\right) + \delta\left(B\Gamma B' + \Lambda\right) \end{aligned} \tag{2.11}$$

Here, B_{SF} denotes the vector of factor loadings for each asset in a restricted single-factor covariance matrix (Σ_{SF}) with factor variance σ^2_{SF} and diagonal matrix of idiosyncratic variances Λ_{SF} and, as before, B, Γ, and Λ represent the corresponding characteristics of an augmented, N factor covariance matrix, and δ represents the shrinkage intensity. The relationship between equations 2.11 and 2.10 is the key to deriving a prior that yields the Ledoit & Wolf estimator as the posterior covariance matrix.

Theorem 2.2 *Suppose the likelihood of the data is given by equation 2.5 and an* *investor's prior belief is given by equation 2.3 with parameters:*

$$\beta_{i,0,k} = \begin{cases} \hat{\beta}_{SF}, & \textit{if } k=1 \\ 0 & \textit{otherwise} \end{cases}, \Omega_{i,0,\{j,k\}} = \begin{cases} \frac{T\delta}{1-\delta}\hat{\sigma}^2_{F_k}, & \textit{if } j=k \\ 0 & \textit{otherwise} \end{cases}$$

Then the posterior covariance matrix is given by the Ledoit & Wolf estimator in equation 2.11.

Proof. The proof for off-diagonal entries in the posterior covariance matrix follows directly from equation 2.10, which simplifies so that the weight assigned to the prior expected factor loadings is constant across factors and assets:

$$\Sigma^* = \delta \sum_{k=1}^{N} \sigma^2_{F_k} \hat{B}_k \hat{B}_k' + (1-\delta) \sum_{k=1}^{N} \sigma^2_{F_k} B_{0,k} B_{0,k}' + \Lambda$$

The proper specifications for s_0 and v_0 will set the matrix $\Lambda = \delta\Lambda_0 + (1-\delta)\hat{\Lambda}$ where Λ_0 is the diagonal matrix with (k,k) entry equal to the idiosyncratic variance estimated in the restricted single factor model and $\hat{\Lambda}$ is the idiosyncratic variance in the unrestricted covariance matrix. This can be done by setting idiosyncratic beliefs so that:

$$v_0 = T\delta, \text{ and, } s_{0,i} = \hat{\sigma}_{\epsilon,i,SF} - \frac{1}{T\delta}\left(\hat{\beta}_i - \beta_{i,0}\right)'\left(\Omega_{i,0} + (F'F)^{-1}\right)^{-1}\left(\hat{\beta}_i - \beta_{i,0}\right)$$

This specification establishes the result:

$$\Sigma^*_{LW} = (1-\delta)\left(\sigma^2_{F_1}\hat{B}_1\hat{B}_1' + \Lambda_0\right) + \delta\left(\sum_{k=1}^{N}\sigma^2_{F_k}\hat{B}_k B_{0,k}' + \hat{\Lambda}\right) = (1-\delta)\Sigma_{SF} + \delta\Sigma_S$$

■

Empirical Bayesian Foundations for General Shrinkage Estimators

The result in the theorem 2.2 immediately extends to shrinkage estimators with any prior factor specification, but sometimes the structured shrinkage target lacks an immediate factor representation. To address this setting, the fact that derived factors are only defined up to scale allows the required flexibility in establishing beliefs consistent with a posterior covariance matrix corresponding to a shrinkage estimator with any positive definite shrinkage target. To begin, denote the shrinkage target using the eigenvalue/eigenvector decomposition for an arbitrary, positive-semidefinite covariance matrix Σ_P as $B_P \Gamma_P B_P'$. Then a shrinkage estimator that shrinks the sample covariance matrix towards Σ_P can be represented as:

$$\Sigma_P^* = (1-\delta)\,\Sigma_P + \delta \Sigma_S = (1-\delta)\,B_P \Gamma_P B_P' + \delta B_S \Gamma_P B_S' + \Lambda_{PS} \tag{2.12}$$

Here, the factors are scaled so that the variance of the kth factor is now equal to the kth eigenvalue of shrinkage target Σ_P. This scaling allows for a uniform shrinkage to apply across all factors.

Theorem 2.3 *Suppose the likelihood of the data is given by equation 2.5 and an investor's prior belief is given by equation 2.3 with parameters:*

$$\beta_{i,0,k} = B_{P\{i,k\}}, \Omega_{i,0,\{j,k\}} = \begin{cases} \frac{T\delta}{1-\delta} \hat{\sigma}_{F_k}^2, & \textit{if } j = k \\ 0 & \textit{otherwise} \end{cases}$$

Then the posterior covariance matrix is given by the shrinkage estimator in equation 2.12.

The proof of theorem 2.3 is almost identical to that for theorem 2.2, and is omitted for brevity. A direct corollary of the above relates to a shrinkage technique proposed by Jagannathan and Ma (2003). They show that non-negativity constraints on the minimum variance portfolio are equivalent to a shrinkage of the covariance matrix determined by the shadow costs of those constraints. In particular, they show that, given covariance matrix Σ_S, a vector shadow costs for each asset's non-negativity constraint λ, and denoting the vector with N ones by 1_N, the constrained minimum variance portfolio is equivalent to the unconstrained minimum variance portfolio for the shrinkage covariance matrix Σ_C^* defined as:

$$\Sigma_C^* = \Sigma_S - 0.5\,(\lambda 1_N' + 1_N \lambda') = 0.5\Sigma_S + 0.5\,(\Sigma_S - \lambda 1_N' - 1_N \lambda') \tag{2.13}$$

Taking the eigenvalue decomposition, define $B_C \Gamma_C B_C' = \Sigma_S - \lambda 1_N' - 1_N \lambda'$ and invoking theorem 2.3 immediately proves the following corollary:

Corollary 2.1 *Suppose the likelihood of the data is given by equation 2.5 and an investor's prior belief is given by equation 2.3 with parameters:*

$$\beta_{i,0,k} = B_{C\{i,k\}}, \Omega_{i,0,\{j,k\}} = \begin{cases} T\hat{\sigma}^2_{F_k}, & \text{if } j = k \\ 0 & \text{otherwise} \end{cases}$$

Then the posterior covariance matrix is given by the Jagannathan & Ma estimator in equation 2.13.

2.3 Finite Sample Expected Loss and Optimal Prior Beliefs

As illustrated in the previous section, given a set of orthogonal factors and a location for prior beliefs about factor loadings, the shrinkage intensity, represented by the parameters $\delta_1, \ldots, \delta_N$ in equation 2.10, corresponds to a set of free parameters for tuning the prior beliefs to optimize a finite sample expected loss function. This section presents the solution to this optimization problem under two specifications for finite sample loss (up to the bandwidth parameter specified in the statistical model). The analysis begins by following Ledoit and Wolf (2003) in presenting optimal priors under the expected Frobenius Norm loss, consistent with an element-by-element mean square error measure for the covariance matrix. Extending this work, I then propose a novel alternative expected loss function focused on addressing covariance matrix estimation for a mean-variance portfolio selection analysis and solve for optimal prior beliefs under this new loss specification.

2.3.1 Optimal Priors for Frobenius Norm Loss (MSE)

To maintain consistency with the existing literature on shrinkage methods, first consider optimal prior beliefs under the the expected Frobenius Loss measure, which also corresponds to the loss function chosen by Ledoit and Wolf (2004a,b) in solving for optimal shrinkage intensities:

$$\mathcal{L} = \left\| \Sigma - \hat{\Sigma} \right\|^2 = \sum_{i=1}^{N} \sum_{j=1}^{N} \left(\sigma_{i,j} - \hat{\sigma}_{i,j} \right)^2 \tag{2.14}$$

The Frobenius loss function is motivated by its relationship to mean-square error and the $\mathcal{L}_2$ norm for matrices, a common loss function for statistical problems. The optimization problem is then to trade off bias and variance from the shrinkage estimator in equation 2.10 to minimize the risk function:

$$\begin{aligned} &\mathcal{R}\left(\delta_1, \delta_2, \ldots, \delta_N\right) \equiv \\ &\qquad E\left[\left\| \Sigma - \sum_{k=1}^{N} \delta_k \sigma_{F_k}^2 \hat{B}_k \hat{B}_k' + \sum_{k=1}^{N} \left(1 - \delta_k\right) \sigma_{F_k}^2 B_{0,k} B_{0,k}' + \Lambda \right\|^2 \right] \end{aligned} \tag{2.15}$$

Squared summations quickly become notationally cumbersome, so denote the total bias and variance for the (i, j) entry of the covariance matrix as:

$$\begin{aligned} \mathcal{B}_{i,j} &= \sum_{q=1}^{N} \sum_{r=1}^{N} \left(\beta_{q,i} \beta_{r,i} - \beta_{0,q,i} \beta_{0,r,i} \right) \left(\beta_{q,j} \beta_{r,j} - \beta_{0,q,j} \beta_{0,r,j} \right) \\ \mathcal{V}_{i,j} &= \sum_{q=1}^{N} \sum_{r=1}^{N} cov \left(\hat{\beta}_{q,i} \hat{\beta}_{r,i}, \hat{\beta}_{q,j} \hat{\beta}_{r,j} \right) \end{aligned}$$

This notation compactly expresses the optimal finite-sample shrinkage intensities (and consequently, the optimal empirical prior beliefs) in the following theorem.

Theorem 2.4 *The risk function in equation 2.15 is minimized when* $\delta_1, \ldots, \delta_N$ *are chosen to equal the solution to the following set of* N *linear equations:*

$$\Psi \delta = \xi \tag{2.16}$$

where:

$$\xi_i = \sum_{q=1}^{N} \sigma_{F,q}^2 \mathcal{B}_{i,q}$$
$$\Psi_{\{i,j\}} = \sigma_{F,j}^2 \left(\mathcal{B}_{i,j} + \mathcal{V}_{i,j}\right)$$

Proof. The mechanical details are somewhat tedious, but they simply involve taking the derivative of the risk function and quite a bit of rudimentary algebra pushing around the orders of summation and simplifying. ■

2.3.2 Optimal Priors for Portfolio Variance Square Error (PVSE)

The second performance statistic I consider is a novel measure based on the squared difference between the estimated variance of a portfolio and the true variance of that portfolio. While this measure is not a common norm for covariance matrices, it is particularly relevant to financial applications and mean-variance analysis. Here, take any absolutely continuous measure over portfolio space, $\Phi(w)$, and integrate over portfolios under this measure to compute an integrated squared difference between the estimated portfolio variance and the true portfolio variance. If the weights are exchangeable under the measure Φ, it is possible to integrate out the portfolio weights and the measure itself from the loss function, defining the loss up to an arbitrary constant of porportionality as:

$$PVSE\left(\Sigma, \hat{\Sigma}\right) = \int \left(w'\Sigma w - w'\hat{\Sigma} w\right)^2 d\Phi\left(w\right) \qquad (2.17)$$
$$\propto \sum_{k=1}^{N}\sum_{l=1}^{N}\sum_{p=1}^{N}\sum_{q=1}^{N} \left(\sigma_{k,l} - \hat{\sigma}_{k,l}\right)\left(\sigma_{p,q} - \hat{\sigma}_{p,q}\right)$$

The key difference between this loss function and the Frobenius norm-based loss function studied by Ledoit & Wolf is that the PVSE measure takes the sum over differences between estimated and true covariances prior to squaring, resulting in a 4-fold summation rather than the double-sum in equation 2.14. This feature captures the tradeoffs in estimating covariances as they relate to the problem of optimal diversification, accounting for the interactions of covariances, rather than the more restrictive Frobenius norm, which only accounts for deviations for individual covariances.

The analysis for deriving optimal priors under the expected PVSE loss is then very similar to that for expected MSE loss in the previous section. However, PVSE loss also accounts for the covariance between errors in estimating the $\{i, j\}$ and $\{p, q\}$ entries in the covariance matrix. As such, define:

$$\begin{aligned}
\mathcal{B}^*_{i,j} =& \sum_{k=1}^{N}\sum_{l=1}^{N}\sum_{p=1}^{N}\sum_{q=1}^{N} \left(\beta_{k,i}\beta_{l,i} - \beta_{0,k,i}\beta_{0,l,i}\right)\left(\beta_{p,j}\beta_{q,j} - \beta_{0,p,j}\beta_{0,q,j}\right) \\
\mathcal{V}^*_{i,j} =& \sum_{k=1}^{N}\sum_{l=1}^{N}\sum_{p=1}^{N}\sum_{q=1}^{N} cov\left(\hat{\beta}_{k,i}\hat{\beta}_{l,i}, \hat{\beta}_{p,j}\hat{\beta}_{q,j}\right)
\end{aligned}$$

As before, this notation allows the optimal finite-sample shrinkage intensities (and consequently, the optimal empirical prior beliefs) to be expressed compactly in the following theorem.

Theorem 2.5 *The risk function in equation 2.17 is minimized when $\delta_1, \ldots, \delta_N$ are chosen to equal the solution to the following set of N linear equations:*

$$\Psi^*\delta = \xi^* \tag{2.18}$$

where:

$$\begin{aligned}
\xi^*_i =& \sum_{q=1}^{N} \sigma^2_{F,q}\mathcal{B}^*_{i,q} \\
\Psi^*_{\{i,j\}} =& \sigma^2_{F,j}\left(\mathcal{B}^*_{i,j} + \mathcal{V}^*_{i,j}\right)
\end{aligned}$$

2.3.3 Feasible Estimation of Optimal Priors

The key to consistently estimating the parameters for optimal priors is to consistently estimate the biases and covariances that collectively define $\mathcal{B}_{i,j}$, $\mathcal{V}_{i,j}$, $\mathcal{B}^*_{i,j}$, and $\mathcal{V}^*_{i,j}$.

This analysis follows the approach of Ledoit and Wolf (2003), who show the bias terms can be consistently estimated by replacing the population moments with unbiased sample moments and taking the difference between the estimated factor loadings and the prior expected factor loadings as follows:

$$\begin{aligned}\hat{\mathcal{B}}_{i,j} =& \sum_{q=1}^{N}\sum_{r=1}^{N}\left(\hat{\beta}_{q,i}\hat{\beta}_{r,i} - \beta_{0,q,i}\beta_{0,r,i}\right)\left(\hat{\beta}_{q,j}\hat{\beta}_{r,j} - \beta_{0,q,j}\beta_{0,r,j}\right) \\ \hat{\mathcal{B}}^*_{i,j} =& \sum_{k=1}^{N}\sum_{l=1}^{N}\sum_{p=1}^{N}\sum_{q=1}^{N}\left(\hat{\beta}_{k,i}\hat{\beta}_{l,i} - \beta_{0,k,i}\beta_{0,l,i}\right)\left(\hat{\beta}_{p,j}\hat{\beta}_{q,j} - \beta_{0,p,j}\beta_{0,q,j}\right)\end{aligned} \tag{2.19}$$

In analyzing the covariance terms, first observe that the orthogonality of the factors immediately implies that $\mathcal{V}^*_{i,j} = \mathcal{V}_{i,j} = 0, \forall i \neq j$ since the covariance between loadings on two different factors will always be zero regardless of the assets. The next step is to obtain a closed form solution for $\mathcal{V}_{i,i}$ and $\mathcal{V}^*_{i,i}$ in terms of population moments that admits a consistent plug-in estimation strategy. After a good deal of algebra to address the N^4 terms in the summand:

$$\begin{aligned}T\mathcal{V}_{i,i} =& 3\sigma_{F_i}^{-4}\sum_{k=1}^{N}\sigma^4_{\epsilon,k} + \sigma_{F_i}^{-4}\sum_{k=1}^{N}\sum_{l\neq k}\sigma^2_{\epsilon,k}\sigma^2_{\epsilon,l} \\ T\mathcal{V}^*_{i,i} =& 3\sigma_{F_i}^{-4}\sum_{k=1}^{N}\sigma^4_{\epsilon,k} + 3\frac{1}{T}\sigma_{F_i}^{-4}\sum_{k=1}^{N}\sum_{l\neq k}\sigma^2_{\epsilon,k}\sigma^2_{\epsilon,l} + 3\sigma_{F_i}^{-2}\sum_{k=1}^{N}\sum_{p\neq k}\sum_{q\neq k,q\neq p}\sigma^2_{\epsilon,k}\beta_{p,i}\beta_{q,i} \\ &+ \sigma_{F_i}^{-2}\sum_{k=1}^{N}\sum_{p\neq k}\sum_{l\neq k,l\neq p}\sigma^2_{\epsilon,p}\beta_{k,i}\beta_{l,i} + \sigma^2_{\epsilon,l}\beta_{k,i}\beta_{p,i}\end{aligned} \tag{2.20}$$

To establish a consistent and feasible estimator of optimal prior beliefs, then, replace the population moments in the equation above with sample moments. The consistency of this estimator for the beliefs follows immediately from the Continuous Mapping Theorem. Consistency of the posterior covariance matrix under optimal priors follows from the fact that the optimal shrinkage places all weight on the sample as the sample estimator becomes arbitrarily precise. The only free parameter remaining to be chosen is the bandwidth parameter h.

2.3.4 Optimal Empirical Bayesian Priors for Structured Estimators

When analyzing priors based on stochastic models that do not have an immediate factor-structure representation, such as the constant-correlation covariance matrix, it would be appropriate to apply the above analysis treating $\beta_{i,0}$ as a random variable. This complication introduces yet another dimension to the optimization problem but does not materially affect the analysis outside of requiring a host of additional algebraic derivation. The main challenges in this setting arise in deriving feasible estimators for optimal priors, as the analysis requires characterizing the covariance of estimators across the models. This exercise is analytically feasible in special cases (such as those explored by Ledoit & Wolf), but it is impossible to derive an analytical result that can be applied to an arbitrary model as an empirical Bayes prior. However, that a sampling or bootstrap approach might provide an generic algorithm that can be generically applied for a fully automated Empirical Bayesian estimator.

2.4 Prior Belief Specifications

Having characterized the posterior covariance matrix's analytical properties and derived finite-sample optimal priors, the discussion now turns to specifications for prior beliefs based on economic research. These beliefs directly relate to factor models inspired by the Bayesian asset pricing literature, highlighting the relationship between the present model and these existing models with two classes of prior beliefs: one that expresses a simple belief that covariances are driven by a k factor structure and a Minnesota-style prior that progressively restricts parameter estimates for non-benchmark factors. In addition to a restrictive model of beliefs that non-benchmark factor loadings are zero, I consider an empirical Bayesian set of beliefs that factor loadings are centered at the cross-sectional mean of OLS-estimated factor loadings. These beliefs capture the empirical regularity of mean-reversion in estimated factor loadings and also address parameter estimation error by shrinking extreme realizations of estimated factor loadings.

2.4.1 No Extra-Benchmark Correlation Prior (BMK)

The introduction proposed a $K < N$ factor model of covariance as the motivating prior belief for the Bayesian Covariance Matrix estimator. In this setting, consider a prior that is diffuse over the first K factor loadings but then shrinks the remaining $N-K$ factor loadings toward zero. As a further simplifying assumption, assume the prior for each factor loading is independent of one another and that the prior standard deviation is constant for each of the remaining $N-K$ factors.

$$\beta_i, \sigma_{\epsilon,i}^2 \sim_{\text{prior}} NG\left(0, \begin{bmatrix} \sigma_\alpha^2 & 0 & 0 \\ 0 & \infty I_K & 0 \\ 0 & 0 & \sigma_C^2 I_{N-K} \end{bmatrix}, v_0, s_0^2\right) \tag{2.21}$$

The approach to formulating this prior is inspired by Pastor (2000) and Pastor and Stambaugh (2002), who model the prior belief in a benchmark asset pricing model with diffuse priors over the factor loadings and informatively shrinks the security's alpha toward zero. Assuming the expected return on the $N-K$ derived factors is zero, the present approach immediately nests the Pastor & Stambaugh model as a special case where $\sigma_{C,0}^2 = 0$.

Under this prior specification, there remains four free parameters to define the prior beliefs. As in the Pastor & Stambaugh pricing model, σ_α characterizes the degree to which the investor believes in the asset pricing model. Similarly, the new parameter σ_C characterizes the degree to which the investor believes in the hypothesis that there is no extra-benchmark correlation across assets. The larger is the value of σ_C, the greater the posterior factor loadings for augmented factors are allowed to deviate from zero. In the extreme case where $\sigma_C \to \infty$, the extra-benchmark factor loadings become freely variable and the posterior covariance matrix converges to the unbiased sample covariance matrix. The prior beliefs for idiosyncratic variance can be set to be diffuse by setting $v_0 = 0$, in which case any finite value for s_0 may be selected without affecting the posterior expectation.

2.4.2 Minnesota Priors (MN)

In the analysis of VAR's, the number of lagged periods included in the analysis presents a significant decision variable in estimating and forecasting in time-series models. Typically, researchers approached the problem using a specification test for the number of lags and dropping all additional lags. The approach to selecting the number of lags closely parallels that used by most researchers who have adopted principal components analysis as a method for dimensionality reduction since Stock and Watson (1989), who select a number of components and restrict the remaining factor loadings to be equal to zero. Connor and Korajczyk (1993) and Bai and Ng (2002) present tests for the number of components to include in the model, but still dogmatically restrict the variation due to other factors.

Using a Bayesian formulation, Litterman (1986) proposed to address the lag selection issue by introducing progessively more restrictive priors as the lag length grows, dubbed the "Minnesota Prior." In the Minnesota prior, the prior beliefs for a parameter associated with the τth lag has expectation equal to zero with a variance inversely proportional to τ. In this way, the Minnesota prior smooths the effect of lag length determinantion as parameters corresponding to longer lags are shrunk towards zero rather than dropped entirely.

A similar approach can be immediately adapted to the number of factors in estimating covariance matrices. Instead of testing for the correct number of factors to include and dropping all additional factors, this approach retains all factors under a prior that the variance for the factor loading is inversely proportional to the index of the factor estimated. Denoting by Ξ_{N-K} the diagonal matrix with entries $\left(1, \frac{1}{2}, \ldots, \frac{1}{N-K}\right)$:

$$\beta_i, \sigma^2_{\epsilon,i} \sim_{\text{prior}} NG\left(0, \begin{bmatrix} \sigma^2_\alpha & 0 & 0 \\ 0 & \infty I_K & 0 \\ 0 & 0 & \sigma^2_C \Xi_{N-K} \end{bmatrix}, v_0, s^2_0\right) \tag{2.22}$$

2.4.3 Mean Reverting Factor Loading Prior (MR)

An alternative prior is motivated by the analysis of Blume (1975) who established the tendency of factor loadings to regress toward the mean. Define the cross-sectional average beta, $\bar{\beta} = \frac{1}{N}\sum_{i=1}^{N} \hat{\beta}_i$ and idiosyncratic variance $\bar{s} = \frac{1}{N}\sum_{i=1}^{N} \hat{s}_i$, so that the investor's prior beliefs shrink the factor loadings toward the grand mean.

$$\beta_i, \sigma^2_{\epsilon,i} \sim_{\text{prior}} NG\left(\bar{\beta}, \begin{bmatrix} \sigma^2_\alpha & 0 & 0 \\ 0 & \infty I_K & 0 \\ 0 & 0 & \sigma^2_C I_{N-K} \end{bmatrix}, v_0, \bar{s}^2_0\right) \tag{2.23}$$

This class of beliefs shrinks factor loadings toward the cross-sectional average factor loading in a manner similar to the asset pricing model proposed by Frost and Savarino (1986). As with the MKT prior, the parameter σ_C represents the degree to which the investor believes in mean-reverting factor loadings. For extremely large values of σ_C, the posterior covariance matrix converges toward the unbiased sample covariance matrix. As $\sigma_C \to 0$, all factor loadings become identical and, in turn, all covariances are shrunk toward a single constant. If v_0 also becomes large, so that all idiosyncratic variances are shrunk towards the same constant, then the posterior covariance matrix converges to a two-parameter covariance matrix with diagonal entries equal to the average variance and off-diagonal entries equal to the average covariance for all assets.

Using the mean-reverting prior introduces an Empirical Bayesian procedure that technically does not satisfy the assumptions of an independent prior and likelihood in the statistical model. However, this prior can be formalized in a learning model where the cross-section is informative about an individual asset's factor loadings akin to that studied in Jones and Shanken (2005) or alternatively in a model with time-varying parameters and mean-reversion in factor loadings. In this sense, the pricing parameter σ_α measures the degree to which an investor believes individual fund alphas can vary from the grand mean alpha, with large values of σ_α allowing individual fund alphas to be effectively unrestricted.

2.5 Monte Carlo Tests & Finite Sample Properties

Having established its statistical properties, the analysis now turns to characterizing the empirical properties of the posterior covariance matrix relative to existing estimators with a simple Monte Carlo test. To maintain consistency with the existing literature on shrinkage methods, first consider the performance of covariance matrix estimators using the Frobenius Loss measure from equation 2.14, though I'll also evaluate estimator performance using the Portfolio Variance Squared Error Loss function proposed in 2.17.

2.5.1 Data Sample and Competing Estimators

In this test, I fit three covariance matrices for $N = 14$, $N = 25$, and $N = 48$ assets, corresponding to the full-sample covariance matrix estimated from 14 country portfolios, 25 Value-Size sorted portfolios, and 48 industry portfolios, respectively, where the return series are downloaded from Ken French's website. For each covariance matrix, I generate a time series of normally distributed random variables where the length of the series (T) ranges from 25 to 500.

I consider two non-Bayesian estimators as reference points: the unrestricted sample covariance matrix and a single-factor model of covariance with the factor derived via principal components. I also include the Ledoit & Wolf shrinkage estimator that shrinks the sample covariance matrix toward the the Single (equal weighted) Factor covariance matrix. I implement the posterior covariance matrix for single-factor prior presented in equation 2.21, the Minnesota prior from equation 2.22, and the mean-reverting prior specification from 2.23 with varying degrees of certainty in the prior specification. In each of the posterior covariance models, I use the normal-inverse gamma regression model with diffuse priors on the idiosyncratic noise, so that the variance of the error term has prior degree of freedoms $v_0 = 0$ and scale parameter $s_0^2 = 0.01$.

2.5.2 Overall Estimator Results

Table 2.1 presents the finite-sample mean square error for each of the estimator at several horizons, with numerical standard errors in parenthesis. Overall, the posterior covariance estimators perform quite well in this exercise, with the mean-reverting prior specification dominating other estimators in the setting where the number of observations is less than the number of securities.

The results in Table 2.1 also illustrate that the posterior covariance estimator mean squared error goes to zero as the number of observations increases. Consequently, as noted in Theorem 2.1, the estimation algorithm retains asymptotic consistency while attaining lower finite-sample mean squared error. This feature is particularly useful when the number of assets in the model is relatively small, as in Panel A of Table 2.1. The relative performance of the estimator is somewhat sensitive to prior beliefs, notably when the prior beliefs are differently centered, as in the case of the Mean-Reverting and Single Factor priors.

The performance of the posterior covariance matrix is less striking when evaluated using the proposed Portfolio Variance Square Error measure in Table 2.2, in large part because this performance measure has so little variation across estimators. Indeed, the sample covariance matrix seems to dominate the other estimation strategies, but in no cases are the differences in performance significantly large.

2.5.3 Individual Estimator Robustness

A common concern with Bayesian estimation techniques is the degree to which estimator results are sensitive to prior belief specifications. This subsection evaluates the effect of varying the degree of belief in each of the underlying factor models using the same simulation test as in the previous subsection. This analysis illustrates the effect of varying the degree of certainty relating to prior beliefs on estimator performance. Due to the lack of broad variation in the PVSE loss measure, I only report results for the MSE, Frobenius Norm, loss measure.

Table 2.3 illustrates the MSE under the Mean-Reverting prior with varying degrees of prior certainty. As the prior standard deviation parameter becomes low, the posterior covariance matrix estimator converges to the model where all factor loadings and idiosyncratic variances are the same across assets. In this case, the assets become exchangeable and the covariance matrix converges to a two parameter model with constant covariances and constant variances. In contrast, as the prior standard deviation becomes large, the posterior covariance matrix converges to the sample covariance matrix.

The performance for the Single-Factor and Minnesota priors are very similar to one another and so table 2.4 reports the loss statistics for the Single Factor prior with varying degrees of variance in the prior beliefs. As the prior beliefs become very restrictive, the posterior expected covariance matrix converges to the single factor covariance matrix estimator. As they become diffuse, the estimator converges to the sample covariance matrix. The best performance levels come with priors that are relatively diffuse compared to the priors under the mean-reverting specification.

Lastly, I turn to the optimal prior specification to evaluate the degree to which the bandwidth parameter can affect performance in this model. Table 2.5 illustrates that for most reasonable bandwidths (1% to 5%), the posterior expected covariance matrix estimated using the optimal prior performs quite well. With extremely small or large bandwidth parameters, however, the estimator's performance begins to deteriorate.

2.6 Performance in Portfolio Selection Exercise

In addition to Monte Carlo exercises characterizing the statistical performance of the posterior covariance matrix, a series of portfolio selection tests characterize the performance of the posterior covariance matrix in an applied setting. To avoid complications due to estimation of means, I focus on the problem of finding global minimum variance portfolio (GMVP) with and without non-negativity constraints. The first test corresponds to a traditional backtest using return histories for the three asset universes analyzed in the Monte Carlo simulation (analysis for many other universes and backtest parameterizations are available upon request) while a second test evaluates the robustness of these results using additional simulation evidence.

2.6.1 Data Sample and Competing Estimators

I evaluate the performance of the Bayesian Factor Covariance estimator in a minimum variance portfolio selection exercise with the same three samples used to calibrate the Monte Carlo tests taken from Ken French's website. On January 1 of each year, I use a rolling window of 10 years' monthly returns to estimate historical covariance matrices and calculate optimal portfolio holdings, rebalancing monthly and reoptimizing annually. The large industry portfolio (corresponding to $N = 48$) and the Country portfolio samples (with $N = 14$) cover only the periods from 1979-2008 and 1975-2008, respectively. The $N = 25$ sample of value and size sorted portfolios has a long history, beginning in 1936.

In this exercise, I include two benchmarks that do not require estimating a covariance matrix. The first is the $1/N$ rule, corresponding to naïve diversification rule of investing an equal proportion of the portfolio in each of the assets in the investible universe. DeMiguel, Garlappi, and Uppal (2007) focused on this strategy and showed it to perform quite well in a variety of empirical applications. The second naïve diversification rule is the $1/V$ rule, where each asset's weight in the portfolio is proportional to the inverse of its variance, which can be thought of arising from a restricted model of zero covariances. It is straightforward to show that, under a diffuse prior for factor loadings, extreme beliefs for the parameters v_0 and s_0 shrink the optimal minimum-variance portfolio weights between these two rules. In addition to the benchmark allocations, the optimal portfolio weights are estimated using the covariance matrix estimators discussed in the previous section.

2.6.2 Minimum-Variance Performance Backtest Results

Tables 2.6 and 2.7 report the performance statistics for unconstrained and constrained minimum-variance portfolios, respectively, based on each of the above estimation strategies. As in the Monte Carlo simulations, the results are strikingly supportive for the Bayesian Factor Covariance estimators. In every sample, portfolios computed using the posterior covariance matrix are among the strategies providing the lowest volatility. The Ledoit & Wolf estimator also does extremely well, illustrating the dominance of Bayesian methods in portfolio selection.

While there is some sensitivity to the prior specification, portfolios derived using the posterior covariance matrix perform remarkably well both in terms of volatility and realized out-of-sample Sharpe Ratios. Interestingly, the $1/N$ strategy is clearly dominated by both the $1/V$ strategy and the Bayesian models in terms of portfolio variance, but still maintains one of the highest Sharpe Ratios and geometric average returns.

2.6.3 Simulated Portfolio Performance Test Results

In addition to the standard backtest analysis characterizing portfolio selection performance in historical samples, I implement the test procedure proposed by Markowitz and Usmen (2003) and also employed by Liechty, Harvey, and Liechty (2008) to evaluate the portfolio selection performance. The Markowitz & Usmen test proceeds in a fashion similar to the Monte Carlo simulation above in that, first, I select a sample of returns from which to compute a true covariance matrix. I then generate a simulated sample of normally distributed returns where the length of the series (T) ranges from 25 to 500. From these samples, I estimate the minimum variance portfolio weights using each of the candidate covariance matrix estimators and compute the true variance of that portfolio for each estimator. I then generate a new sample and repeat the exercise until the numerical standard error for the volatility of the GMVP generated by each estimation strategy below one basis point (which takes less than 1,000 iterations).

Table 2.8 presents the average volatility attained by each of the estimators in identifying the Global Minimum Variance Portfolio with no long/short constraints. The Bayesian estimators perform quite well relative to both the naïve portfolio strategies and the optimal weights estimated using non-Bayesian methods. The settings where the Bayesian estimators underperform correspond to cases with a large number of observations relative to the number of securities. These results are still somewhat sensative to the prior belief specification and especially so in other specifications where the prior beliefs are extremely tight. However, the portfolio weights estimated using moderate prior beliefs perform well across most of the specifications.

The performance of the covariance matrix estimators in the constrained portfolio selection problem is presented in table 2.9. As in the backtest study, the constraints reduce the influence of prior beliefs so that the realized GMVP volatility from the posterior covariance matrices is always within 20 basis points of the best alternative strategy even in very small samples. Interestingly, the constraints substantially improve the performance of portfolios estimated using the sample covariance matrix and the single factor covariance matrix estimators, particularly when the sample size is relatively small. However, these constraints do little to improve the performance of portfolios calculated using the Bayesian covariance matrix estimators. Indeed, under most prior specifications and sample sizes, the constrained portfolios have higher volatility than the unconstrained portfolios.

2.7 Conclusion and Potential Extensions

I present a formal foundation for applying Bayesian methods to estimating covariance matrices, showing that my analysis theoretically nests a broad set of popular Bayesian techniques for estimating means and covariances. The Monte

Carlo and minimum variance portfolio tests are quite encouraging for the prospective application of the posterior covariance matrix estimator. Another potentially valuable application would include adapting these methods to modeling the square error of different forecast errors in the forecast combination literature.

One limitation of the current exercise is that it's inherently static in its consideration of the data generating process. However, much of the reasoning behind mean-reverting factor loadings is based on a dynamic specification of the return generating process with time-varying factor loadings. The current model could readily be extended to a conditional setting where factor loadings follow a mean-reverting AR(1) process, with the posterior estimates of the persistence characterizing the degree of mean reversion. This model would also provide a purely Bayesian framework to replace the empirical Bayesian framework used above with likely similar results.

Another open question relates to the underlying mechanisms by which the Bayesian Factor Covariance estimator addresses parameter and model uncertainty. One appeal of the estimator is that it allows the data to dictate the structure of the model. Introducing a cross-validation step in prior selection to specify the bandwidth parameter could deliver a fully automated estimator with quite appealing out of sample properties. This estimator would be very easy to implement and, as such, likely be quite useful in a variety of economic settings.

Tables and Figures

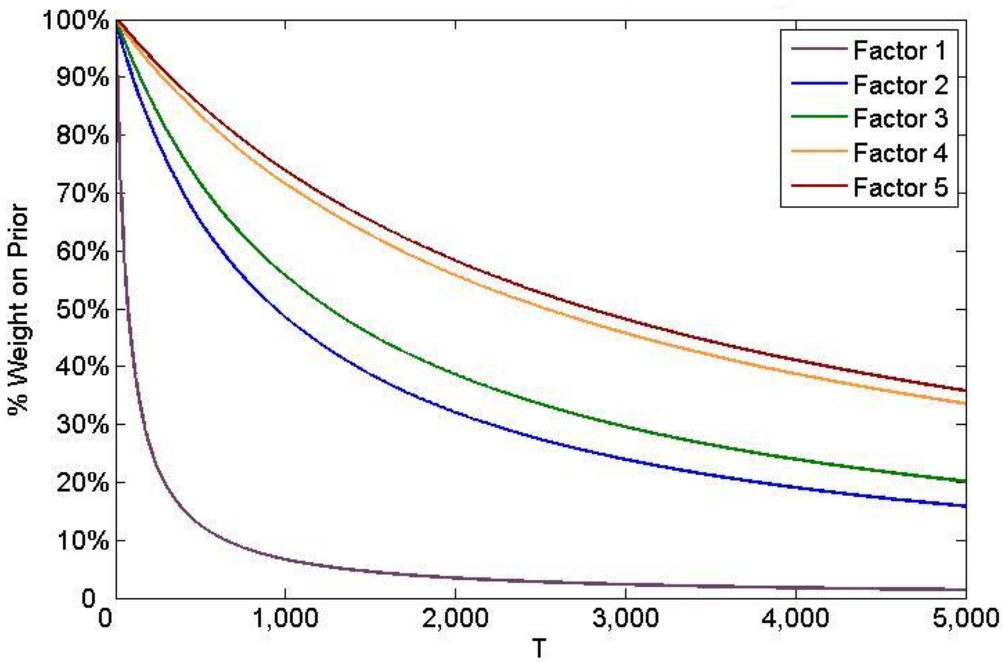

Figure 2.1: Posterior Factor Loading Consistency
As the number of observations becomes large, the weight assigned to prior factor loadings goes to zero. In this way, the posterior expected covariance matrix converges to the sample covariance matrix, which is consistent for the true covariance matrix.

Table 2.1: Monte Carlo Simulation Mean Squared Error
This table presents the Mean-Squared Error in estimating full-sample covariance matrix from normally distributed simulation data of various horizons. These results are based on 10,000 simulations. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5.

		Single	Market	Bayesian Factor Covariance Matrices			
T	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
			Panel A: 14 Country Equity Portfolios				
25	19.13	19.85	17.16	12.18	18.70	18.77	13.30
50	9.41	11.07	8.74	7.84	9.50	9.70	8.61
75	6.21	8.17	5.87	6.10	6.33	6.59	6.44
100	4.60	6.73	4.39	5.09	4.71	4.97	5.03
250	1.84	4.29	1.81	2.80	1.88	2.03	2.04
500	0.91	3.46	0.90	1.61	0.93	0.99	1.00
			Panel B: 25 Size & Value Sorted Portfolios				
25	160.02	162.91	159.18	155.63	160.03	160.58	156.34
50	79.68	83.29	79.44	80.72	79.80	80.23	80.45
75	52.56	56.43	52.45	54.53	52.65	52.96	54.78
100	39.41	43.37	39.34	41.63	39.48	39.70	42.32
250	15.67	19.83	15.66	17.37	15.69	15.77	18.49
500	7.65	11.85	7.65	8.56	7.66	7.69	9.66
			Panel C: 48 Industry Portfolios				
25	235.70	226.04	213.78	185.00	226.14	223.78	181.65
50	112.48	116.20	104.21	100.81	109.89	109.15	102.56
75	75.52	83.71	71.09	72.85	74.50	74.25	76.59
100	57.40	67.87	54.64	57.98	56.93	56.92	62.43
250	22.62	37.41	22.10	25.37	22.64	22.89	28.23
500	11.45	27.61	11.31	13.18	11.47	11.64	13.75

Table 2.2: Monte Carlo Simulation Portfolio Variance Squared Error
This table presents the Portfolio Variance Squared Error (defined in equation 2.17) in estimating full-sample covariance matrix from normally distributed simulation data of various horizons. These results are based on 10,000 simulations. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5.

		Single	Market	Bayesian Factor Covariance Matrices			
T	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
			Panel A: 14 Country Equity Portfolios				
25	13.74	14.77	13.99	14.43	13.95	14.08	14.66
50	6.80	7.71	6.97	7.48	6.87	6.97	7.76
75	4.48	5.33	4.58	5.05	4.51	4.58	5.34
100	3.27	4.12	3.34	3.79	3.29	3.34	4.08
250	1.32	2.19	1.34	1.63	1.33	1.35	1.78
500	0.65	1.51	0.66	0.81	0.65	0.66	0.80
			Panel B: 25 Size & Value Sorted Portfolios				
25	75.61	75.71	75.63	75.66	75.64	75.65	75.71
50	37.80	37.88	37.81	37.87	37.82	37.83	37.93
75	24.92	24.99	24.93	24.98	24.93	24.94	25.03
100	18.70	18.75	18.70	18.74	18.71	18.71	18.79
250	7.45	7.49	7.45	7.47	7.45	7.45	7.49
500	3.62	3.65	3.62	3.62	3.62	3.62	3.62
			Panel C: 48 Industry Portfolios				
25	23.94	24.03	23.92	23.92	23.95	23.96	23.92
50	10.97	11.01	10.97	10.98	10.97	10.97	11.00
75	7.48	7.50	7.47	7.48	7.47	7.47	7.50
100	5.79	5.82	5.78	5.80	5.79	5.79	5.82
250	2.22	2.28	2.22	2.24	2.22	2.22	2.28
500	1.15	1.22	1.15	1.16	1.15	1.15	1.19

Table 2.3: Mean Reverting Prior Mean Squared Error Performance
This table presents the Mean-Squared Error in estimating full-sample covariance matrix from normally distributed simulation data of various horizons based on 10,000 simulations using the Bayesian Factor Covariance Matrix with a Mean Reverting prior and varying levels of prior certaintly. As the prior standard deviation becomes large, the estimator converges to the sample estimator. As the prior standard deviation becomes small, the estimator converges to a two parameter covariance matrix with constant covariances and constant variances.

	Mean Reverting Prior with Standard Deviation Parameter:							
T	0.01	0.1	0.5	1	2.5	5	10	Sample
	Panel A: 14 Country Equity Portfolios							
25	11.87	11.84	11.48	12.18	15.46	17.46	18.70	19.13
50	8.03	7.97	7.39	7.84	8.74	9.16	9.37	9.41
75	6.70	6.61	5.89	6.10	6.14	6.18	6.21	6.21
100	6.03	5.92	5.08	5.09	4.72	4.61	4.60	4.60
250	4.94	4.67	3.46	2.80	2.00	1.86	1.84	1.84
500	4.55	4.06	2.55	1.61	0.99	0.92	0.91	0.91
	Panel B: 25 Size & Value Sorted Portfolios							
25	240.99	233.82	170.48	155.63	158.31	159.07	159.62	160.02
50	179.93	166.49	89.62	80.72	79.95	79.54	79.61	79.68
75	159.05	139.86	60.21	54.53	52.97	52.53	52.54	52.56
100	148.88	124.50	45.76	41.63	39.77	39.40	39.40	39.41
250	130.26	82.07	19.66	17.37	15.81	15.68	15.67	15.67
500	123.32	52.60	10.65	8.56	7.70	7.65	7.65	7.65
	Panel C: 48 Industry Portfolios							
25	215.91	210.77	174.40	185.00	208.62	222.55	231.36	235.70
50	158.48	148.90	99.62	100.81	104.44	108.93	111.57	112.48
75	142.84	129.15	75.06	72.85	71.88	74.05	75.22	75.52
100	135.27	117.76	61.98	57.98	55.46	56.67	57.28	57.40
250	119.27	83.82	32.38	25.37	22.57	22.59	22.62	22.62
500	114.00	60.79	19.47	13.18	11.53	11.45	11.45	11.45

Table 2.4: Single Factor Prior Mean Squared Error Performance

This table presents the Mean-Squared Error in estimating full-sample covariance matrix from normally distributed simulation data of various horizons based on 10,000 simulations using the Bayesian Factor Covariance Matrix with a Mean Reverting prior and varying levels of prior certaintly. As the prior standard deviation becomes large, the estimator converges to the sample estimator. As the prior standard deviation becomes small, the estimator converges to a two parameter covariance matrix with constant covariances and constant variances.

	Single	Single Factor Prior with Standard Deviation Parameter:							
T	Factor	0.01	0.1	0.5	1	2.5	5	10	Sample
Panel A: 14 Country Equity Portfolios									
25	19.85	19.83	19.79	19.65	19.10	18.85	18.70	19.08	19.13
50	11.07	11.06	11.02	10.77	10.00	9.76	9.50	9.47	9.41
100	6.73	6.73	6.64	6.16	5.23	5.02	4.71	4.63	4.60
250	4.29	4.29	3.99	3.07	2.21	2.06	1.88	1.85	1.84
500	3.46	3.45	2.78	1.70	1.08	1.00	0.93	0.92	0.91
Panel B: 25 Size & Value Sorted Portfolios									
25	162.91	162.89	162.86	162.69	161.79	161.27	160.03	159.91	160.02
50	83.29	83.29	83.24	82.88	81.40	80.80	79.80	79.69	79.68
100	43.37	43.37	43.24	42.42	40.50	40.03	39.48	39.42	39.41
250	19.83	19.83	19.35	17.61	16.05	15.86	15.69	15.67	15.67
500	11.85	11.85	10.69	8.64	7.78	7.72	7.66	7.65	7.65
Panel C: 48 Industry Portfolios									
25	226.04	225.93	225.55	223.79	220.07	220.10	226.14	232.38	235.70
50	116.20	116.17	115.56	112.48	107.82	107.57	109.89	111.84	112.48
100	67.87	67.85	66.54	61.35	56.62	56.33	56.93	57.35	57.40
250	37.41	37.39	33.29	26.04	22.93	22.73	22.64	22.63	22.62
500	27.61	27.55	19.72	13.37	11.69	11.57	11.47	11.45	11.45

Table 2.5: Optimal Prior Mean Squared Error Performance

This table presents the Mean-Squared Error in estimating full-sample covariance matrix from normally distributed simulation data of various horizons based on 10,000 simulations. The optimal prior covariance matrix performs rather well for a broad set of moderate bandwidth parameters. However, as the bandwidth becomes extremely small or extremely large, performance begins to deteriorate.

	Optimal Prior with Bandwidth Parameter:								
T	0.005	0.01	0.015	0.025	0.05	0.075	0.1	0.25	Sample
Panel A: 14 Country Equity Portfolios									
25	16.39	15.56	15.35	14.84	13.30	12.66	12.48	12.59	19.13
50	8.77	8.45	8.34	8.55	8.61	8.22	7.94	7.90	9.41
100	4.35	4.42	4.39	4.42	5.03	5.37	5.41	5.25	4.60
250	3.77	2.02	1.83	1.87	2.04	2.34	2.70	3.59	1.84
500	4.66	3.41	1.93	1.03	1.00	1.10	1.23	2.33	0.91
Panel B: 25 Size & Value Sorted Portfolios									
25	151.75	153.00	153.53	154.07	156.34	160.88	167.11	202.48	160.02
50	76.89	77.83	78.75	79.59	80.45	81.70	83.66	106.74	79.68
100	39.56	39.37	39.86	41.08	42.32	42.77	43.23	49.94	39.41
250	17.90	17.49	17.25	17.20	18.49	19.52	20.02	21.22	15.67
500	10.51	10.35	10.11	9.80	9.66	10.25	10.99	12.79	7.65
Panel C: 48 Industry Portfolios									
25	210.16	207.80	205.66	197.27	181.65	180.42	184.98	206.26	235.70
50	101.63	101.11	102.80	104.67	102.56	101.31	103.24	130.32	112.48
100	54.26	53.95	55.03	58.27	62.43	63.00	63.23	77.63	57.40
250	22.28	22.08	22.27	23.52	28.23	31.94	34.03	37.63	22.62
500	11.60	11.54	11.50	11.76	13.75	16.47	19.10	26.34	11.45

Table 2.6: Performance in Portfolio Backtest Exercise

This table presents the out of sample portfolio performance for annually rebalanced portfolios whose weights are calculated to minimize variance using the estimated covariance matrix from a sample of 120 lagged returns, using data provided by Ken French. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5. Notably, the Bayesian estimators perform quite well compared to both naïve, sample, and the single-factor estimators.

				Single	Market	Bayesian Factor Covariance Matrices			
	1/N	1/V	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
Panel A: 14 Country Equity Portfolios									
Volatility	15.98%	15.75%	16.32%	16.94%	16.21%	16.45%	16.37%	16.52%	16.06%
Geometric Mean	16.58%	16.84%	18.40%	17.78%	18.30%	17.74%	18.46%	18.43%	17.01%
Arithmetic Mean	16.74%	16.92%	18.35%	17.92%	18.25%	17.80%	18.41%	18.41%	17.13%
Sharpe Ratio	0.735	0.757	0.818	0.763	0.818	0.778	0.819	0.812	0.755
Panel B: 25 Size & Value Sorted Portfolios									
Volatility	19.02%	17.57%	13.50%	17.85%	13.48%	16.70%	13.46%	13.63%	13.89%
Geometric Mean	12.88%	12.93%	12.80%	10.65%	12.59%	11.47%	12.15%	11.94%	12.24%
Arithmetic Mean	14.00%	13.78%	13.02%	11.78%	12.83%	12.31%	12.43%	12.26%	12.56%
Sharpe Ratio	0.473	0.500	0.594	0.380	0.581	0.438	0.552	0.533	0.545
Panel C: 48 Industry Portfolios									
Volatility	15.43%	14.71%	12.51%	13.03%	11.83%	12.17%	12.44%	12.19%	11.98%
Geometric Mean	13.17%	13.30%	13.60%	14.41%	13.73%	13.89%	13.47%	13.97%	14.34%
Arithmetic Mean	13.65%	13.65%	13.60%	14.38%	13.63%	13.80%	13.47%	13.89%	14.19%
Sharpe Ratio	0.560	0.588	0.687	0.720	0.730	0.724	0.681	0.729	0.767

Table 2.7: Performance in Constrained Portfolio Backtest Exercise

This table presents the out of sample portfolio performance for annually rebalanced portfolios whose weights are calculated to minimize variance using the estimated covariance matrix from a sample of 120 lagged returns, using data provided by Ken French. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5.

				Single	Market	Bayesian Factor Covariance Matrices			
	1/N	1/V	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
Panel A: 14 Country Equity Portfolios with Non-Negativity Constraint									
Volatility	15.98%	15.75%	15.93%	15.92%	15.86%	15.84%	15.97%	16.00%	15.80%
Geometric Mean	16.58%	16.84%	17.79%	17.81%	17.87%	17.87%	17.90%	18.03%	18.08%
Arithmetic Mean	16.74%	16.92%	17.77%	17.78%	17.83%	17.82%	17.87%	17.99%	17.99%
Sharpe Ratio	0.735	0.757	0.802	0.803	0.809	0.809	0.806	0.812	0.822
Panel B: 25 Size & Value Sorted Portfolios with Non-Negativity Constraint									
Volatility	19.02%	17.57%	14.75%	14.79%	14.75%	14.77%	14.72%	14.70%	14.78%
Geometric Mean	12.88%	12.93%	11.38%	11.26%	11.37%	11.16%	11.31%	11.30%	11.17%
Arithmetic Mean	14.00%	13.78%	11.92%	11.81%	11.91%	11.72%	11.85%	11.84%	11.73%
Sharpe Ratio	0.473	0.500	0.469	0.461	0.469	0.455	0.465	0.465	0.455
Panel C: 48 Industry Portfolios with Non-Negativity Constraint									
Volatility	15.43%	14.71%	12.22%	12.03%	12.07%	12.06%	12.22%	12.17%	12.00%
Geometric Mean	13.17%	13.30%	12.66%	14.12%	12.74%	13.47%	12.71%	12.60%	13.82%
Arithmetic Mean	13.65%	13.65%	12.72%	14.00%	12.78%	13.43%	12.76%	12.66%	13.73%
Sharpe Ratio	0.560	0.588	0.632	0.748	0.644	0.699	0.635	0.629	0.727

Table 2.8: Performance in Unconstrained Portfolio Selection Simulation

This table presents the true realized portfolio volatility for portfolios whose weights are calculated to minimize variance using the estimated covariance matrix from a sample of T normally distributed observed returns. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5.

				Single	Market	Bayesian Factor Covariance Matrices			
T	1/N	1/V	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
Panel A: 14 Country Equity Portfolios									
25	4.56%	4.49%	5.99%	5.00%	4.94%	4.63%	4.97%	4.93%	4.78%
50	4.56%	4.48%	4.92%	4.84%	4.68%	4.60%	4.82%	4.80%	4.68%
100	4.56%	4.48%	4.54%	4.72%	4.49%	4.54%	4.58%	4.63%	4.54%
250	4.56%	4.48%	4.36%	4.62%	4.35%	4.42%	4.38%	4.42%	4.37%
500	4.56%	4.48%	4.30%	4.58%	4.30%	4.35%	4.31%	4.32%	4.31%
Panel B: 25 Size & Value Sorted Portfolios									
25	6.95%	6.46%	5.55%	5.54%	5.14%	5.70%	5.15%	5.24%	5.34%
50	6.95%	6.46%	4.92%	5.58%	4.85%	5.67%	4.88%	5.02%	5.00%
100	6.95%	6.46%	4.61%	5.60%	4.61%	5.52%	4.64%	4.71%	4.69%
250	6.95%	6.45%	4.44%	5.60%	4.44%	5.15%	4.45%	4.50%	4.79%
500	6.95%	6.46%	4.37%	5.60%	4.37%	4.71%	4.38%	4.41%	4.66%
Panel C: 48 Industry Portfolios									
25	4.80%	4.60%	4.47%	3.97%	3.90%	4.02%	3.98%	3.94%	3.92%
50	4.80%	4.60%	4.24%	3.88%	3.66%	3.83%	3.77%	3.72%	3.72%
100	4.80%	4.60%	3.69%	3.89%	3.48%	3.67%	3.59%	3.56%	3.55%
250	4.80%	4.60%	3.32%	3.91%	3.30%	3.45%	3.34%	3.36%	3.34%
500	4.80%	4.60%	3.21%	3.92%	3.21%	3.33%	3.22%	3.24%	3.23%

Table 2.9: Performance in Constrained Portfolio Selection Simulation

This table presents the true realized portfolio volatility for portfolios whose weights are calculated to minimize variance using the estimated covariance matrix from a sample of T normally distributed observed returns. The Bayesian Factor Covariance Matrix is estimated using Mean Reverting (MR) priors with unit prior standard error parameter, Single Factor (SF) priors with standard error parameter 5, Minnesota (MN) priors with standard error parameter 5, and the Optimal prior with bandwidth parameter 0.5. Here, portfolio weights are constrained to be non-negative.

				Single	Market	Bayesian Factor Covariance Matrices			
T	1/N	1/V	Sample	Factor	Shrinkage	MR-1	SF-5	MN-5	Opt
Panel A: 14 Country Equity Portfolios									
25	4.56%	4.49%	4.68%	4.64%	4.64%	4.51%	4.65%	4.65%	4.59%
50	4.56%	4.48%	4.52%	4.52%	4.50%	4.47%	4.53%	4.53%	4.50%
100	4.56%	4.48%	4.42%	4.44%	4.41%	4.42%	4.43%	4.44%	4.43%
250	4.56%	4.48%	4.34%	4.38%	4.33%	4.35%	4.34%	4.35%	4.34%
500	4.56%	4.48%	4.30%	4.35%	4.30%	4.31%	4.30%	4.31%	4.30%
Panel B: 25 Size & Value Sorted Portfolios									
25	6.95%	6.46%	5.34%	5.26%	5.30%	5.30%	5.30%	5.29%	5.26%
50	6.95%	6.46%	5.24%	5.20%	5.22%	5.21%	5.22%	5.21%	5.20%
100	6.95%	6.45%	5.18%	5.17%	5.17%	5.17%	5.17%	5.17%	5.17%
250	6.95%	6.46%	5.16%	5.16%	5.15%	5.15%	5.16%	5.15%	5.15%
500	6.95%	6.46%	5.15%	5.16%	5.15%	5.15%	5.15%	5.15%	5.15%
Panel C: 48 Industry Portfolios									
25	4.80%	4.60%	4.06%	3.94%	3.96%	3.95%	4.00%	3.98%	3.95%
50	4.80%	4.60%	3.83%	3.78%	3.78%	3.78%	3.81%	3.79%	3.78%
100	4.80%	4.60%	3.70%	3.71%	3.68%	3.69%	3.70%	3.69%	3.70%
250	4.80%	4.60%	3.62%	3.66%	3.62%	3.64%	3.62%	3.63%	3.64%
500	4.80%	4.60%	3.60%	3.65%	3.59%	3.61%	3.60%	3.60%	3.62%

Bibliography

Jushan Bai and Serena Ng. Determining the number of factors in approximate factor models. Econometrica, 70(1):191–221, January 2002. URL http://ideas.repec.org/a/ecm/emetrp/v70y2002i1p191-221.html.

ME Blume. Betas and Their Regression Tendencies. Journal of Finance, 30(3): 785–795, 1975. ISSN 0022-1082.

Gregory Connor and Robert A Korajczyk. A test for the number of factors in an approximate factor model. Journal of Finance, 48(4):1263–91, September 1993. URL http://ideas.repec.org/a/bla/jfinan/v48y1993i4p1263-91.html.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? Review Financial Studies, page hhm075, 2007. doi: 10.1093/rfs/hhm075. URL http://rfs.oxfordjournals.org/cgi/content/abstract/hhm075v1.

PA Frost and JE Savarino. An Empirical Bayes Approach to Efficient Portfolio Selection. Journal of Financial and Quantitative Analysis, 21(3):293–305, SEP 1986. ISSN 0022-1090.

John Geweke. Contemporary Bayesian Econometrics and Statistics. Wiley-Interscience, 2005.

R Jagannathan and TS Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. Journal of Finance, 58(4):1651–1683, AUG 2003. ISSN 0022-1082.

CS Jones and J Shanken. Mutual fund performance with learning across funds. Journal of Financial Economics, 78(3):507–552, DEC 2005. ISSN 0304-405X. doi: 10.1016/j.jfineco.2004.08.009.

Gary Koop. Bayesian Econometrics. Wiley-Interscience, 2003.

O Ledoit and M Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. Journal of Empirical Finance, 10:603–621, 2003. ISSN 0047-259X.

O Ledoit and M Wolf. A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis, 88(2):365–411, FEB 2004a. ISSN 0047-259X. doi: 10.1016/S0047-259X(03)00096-4.

O Ledoit and M Wolf. Honey, I shrunk the sample covariance matrix - Problems in mean-variance optimization. Journal of Portfolio Management, 30(4):110+, SUM 2004b. ISSN 0095-4918.

John Liechty, Campell Harvey, and Merrill Liechty. Bayes vs. resampling: A rematch. Journal of Investment Management, 6:1–17, 2008.

Robert B Litterman. Forecasting with bayesian vector autoregressions-five years of experience. Journal of Business & Economic Statistics, 4(1):25–38, January 1986. URL `http://ideas.repec.org/a/bes/jnlbes/v4y1986i1p25-38.html`.

Harry Markowitz and Nilufer Usmen. Resampled frontiers vs diffuse bayes: An experiment. Journal Of Investment Management, 1:9–25, 2003.

L Pastor. Portfolio selection and asset pricing models. Journal of Finance, 55(1): 179–223, FEB 2000. ISSN 0022-1082.

L Pastor and RF Stambaugh. Investing in equity mutual funds. Journal of Financial Economics, 63(3):351–380, MAR 2002. ISSN 0304-405X.

James H. Stock and Mark W. Watson. New indexes of coincident and leading economic indicators. NBER Macroeconomics Annual, 4:351–394, 1989. ISSN 08893365. URL `http://www.jstor.org/stable/3584985`.

Chapter 3

An Extended Theory of Conformity

This paper analyzes conformist tendencies for a population in which individuals gain utility by mimicking the average behavior, characterizing norms by the mean behavior. In so doing, the model extends Bernheim's "A Theory of Conformity" (1994) by introducing an endogenous mechanism for establishing social norms. The most interesting result is that this extension does not alter the properties of equilibria established in Bernheim's initial development, that is, social preferences generally give rise to more concentrated behavior and a conformist pool forms when social preferences are sufficiently prominent. Further the extension introduces no new equilibria, since even though Bernheim's development included a multiplicity of locations for conformist outcomes, these outcomes are identified exactly by the location of the social norm within the extended model. In addition to illustrating the determinants of conformist behavior with an endogenous reference point, these findings support applied work inferring social norms from average behavior.

3.1 Introduction

Convergent and conformist behavior arises in a wide variety of economic contexts. In financial markets, traders herding together often have profound effects that are associated with market booms (such as the dot-com era and the Dutch tulip mania) and crashes (from the 1929 crash to the dot-com bust), or coordinated currency attacks that force a government's monetary policy into alignment with market fundamentals. Beyond the context of financial markets, conformity and convergent behavior is observed in a wide variety of everyday economic interactions. From the establishment of social and cultural norms such as tipping to the development of fashions and fads dictating consumer purchases, economic behavior is replete with occasions where individual behavior is in perfect or near perfect concert with the actions of other agents.

An extensive literature analyzing herd behavior has developed under the rubric of coordination games, with the earliest work focusing on coordination in general and considering how agents strategically coordinate. Commonly, discrete coordination games have multiple equilibria that can be characterized as Pareto optimal coordination, inefficient coordination, and failed coordination with random matching.[1] A continuous coordination game setting with these outcomes is related to the study of network effects, the roots of which trace back to Rohlfs (1974). In network effects games, an agents utility from a product is increasing in the user base of that product, and often focuses on conditions that give rise to suboptimal coordination on technological standards and the impact of network effects on efficiency. The outcome of convergent behavior in models of coordination games and network effects is not surprising as the strategic settings are constructed to model coordinated outcomes. While these games properly represent the phenomenon of coordination, they do not typically capture the true nature of conformity, which corresponds to a spontaneous coordination across individuals despite heterogeneous preferences. Modeling this type of behavior requires constructing a setting where coordination is not pre-ordained in order to identify conditions under which coordination may arise naturally.

The game theory literature contains two prominent models for attaining conformist behavioral outcomes in settings with diverse agents and preferences. Among the most successful models in application is the model of information cas-

[1] See, for example, Fudenberg and Tirole (1991).

cades developed by Banerjee (1992) and Bikhchandani, Hirschleifer, and Welch (1992). In an information cascade, conformity takes the form of censored information and herd actions driven by information aggregation. The model consists of a fully observable sequential decision structure where information revealed through the history of decisions eventually overwhelms any individual's information. This remarkably powerful model has been applied in numerous settings, notably the modeling of investor behaviors that generate market booms and busts, sequential voting modeling, and verified in laboratory settings.[2] Similar types of cascade models, for example applications of the Ising model of particle charge and rotation from physics, have been used to model herd behavior and motivate fat tails or return distributions observed in the empirical finance literature (see, for example, Cont and Bouchaud (2000) and Chowdhury and Stauffer (1999)).

A completely continuous, preference based approach to conformity was initially developed in Bernheim (1994)'s "A Theory of Conformity". In this model, Bernheim introduces a signaling game where individuals have an intrinsic incentive to reveal their type but also a socially motivated incentive to be perceived as a particular, almost surely different, type. The reduced form of social preferences adopted by Bernheim directly incorporates social preferences into an individual's utility function and accommodates a broad set of motivations, including settings where conformity might arise from peer effects similar to those in Akerlof (1980), or due to a post game event such as the according of status (for example, tipping or fashion fads), or the likelihood of an investigation (as may be the case when trading stocks based on insider information).[3]

[2]There is a tremendous literature citing the initial Bikhchandani, Hirschleifer, and Welch (1992) and Banerjee (1992) papers. Each of the founding authors of this work have written papers on the implications of cascades for financial markets, for example, see Bikhchandani and Sharma (2001), Devenow and Welch (1996), Hirshleifer and Teoh (2001). The application of information cascades to sequential voting is a fairly new area of research, including the work by Ali and Kartik (2006), Battaglini (2005) and Dekel and Piccione (2000). For experimental work, note Anderson and Holt (1997) and Hung and Dominitz (2004).

[3]Another model of social preferences, where every agent has utility that is monotonic in their rank, is considered in a series of papers by Cole, Mailath, and Postlewaite (1992, 2001) and advocated by Postlewaite (1998). While the current analysis could be adapted to accommodate such preferences, these preferences do not satisfy the sensitivity required to establish Theorems 4 and 5, below, that further characterize the conformity equilibrium. In particular, their arguments are not dense in the players' action space.

As a motivating example for the current model, suppose an individual's type revealed to them insider information about a major corporation's earnings. If markets are assumed to be efficient, then the average investor would have no information beyond what is incorporated in the corporation's stock price and maintain a neutral (market) investment position in the asset. In contrast, an individual whose type reveals significant insider information could maximize their individual wealth by taking an extreme trading position to exploit that information (for example using short sales or option purchases). In a post game stage, suppose the Securities Exchange Commission chooses whether or not to investigate each investor based exclusively on their publicly observed investment decision. Suppose further that an investigation is more likely for individuals who take more extreme trading positions. This threat of investigation gives the informed individual an incentive to appear as if they were any other investor, even at the cost of maximizing individual wealth by fully exploiting their private information. If the cost of being punished were considered high relative to the potential benefits from trading, the investor may choose not to take advantage of the information at all and simply herd with the market position.

This paper addresses the above insider trading example and similar settings by extending the Bernheim approach to model games where the individual's social preferences are dictated by the normative behavior. The original development is not trivially applicable to this setting, as the normative (market) behavior is not exogenously defined but rather depends on the actions of all agents. From the insider trading example, if every agent expects low earnings, that expectation will already be incorporated into the price and the relative value of the individual's information (the "social preference intensity" defined later) will be lower. In particular, the normative behavior is defined here as the population expected action, and can be further extended to include a variety of mechanisms that establish social norms in the form of an optimal type that players wish to be perceived as. The particular case where social norms correspond to the expected action is motivated by recent experimental studies, including Andreoni and Bernheim (2009)'s theoretical and experimental analysis on identifying normative behavior with social preferences.

The analysis proceeds as follows, Section 3.2 formally defines the model by laying out preferences, decisions, and beliefs for players in the game, with the equilibrium concept characterized in Section 3.3. Section 3.4 establishes the necessary and sufficient conditions for a fully separating equilibrium to exist and the equilibria with incomplete separation are treated in Section 3.5. Section 3.6 further extends the analysis by considering several cases of non-trivial social preference intensities and presents some comparative statics results in that setting. Section 3.7 comments on other, potentially interesting questions related to the current model and Section 3.8 concludes.

3.2 The Model for Preferences and Actions

The setting for the current model is identical to that of Bernheim with some minor changes to notation to adopt consistency with the broader signaling literature. Here assume there is a large number, I, of individual agents, indexed by i, who are each privately assigned a type $t_i \in [0, 2] \equiv T$. Players do not directly observe each other's types, but each player chooses a public signal, $a_i \in [0, 2] \equiv A$, that other players observe and use to infer individual i's true type. The true types are assumed to be independently drawn from the set of types according to a cumulative distribution function $F(\cdot)$, with corresponding probability density function $f(\cdot)$, where the cumulative density assumed to be strictly increasing with $F(2) = 1$.

The Bernheim model specifies preferences represented by a utility function with two components, the first representing *intrinsic* utility and the second representing *social* utility. Intrinsic utility measures the extent to which an individual's action deviates from their actual type and is represented by the intrinsic utility function $g(a - t)$, which is assumed to be twice continuously differentiable, strictly concave, symmetric, and maximized at $a = t$.[4] The interpretation of this specifi-

[4]As noted by Bernheim, the concavity and symmetry are assumed primarily for convenience where as the differentiability requirements are imposed on the model to ensure that the conformity result is not driven primarily by some structural discontinuity.

cation is that the individual's type represents their "Intrinsic Bliss Point" (IBP). From the insider trading example above, the IBP corresponds to the optimal asset position for a player based on their information and absent any considerations of potential prosecution. Similarly, the intrinsic utility represents the amount of wealth the player believes they would attain for any given market position given their information.

In addition to the intrinsic utility component of preferences, the Bernheim model incorporates a social utility component with properties similar to intrinsic utility, with social utility maximized when an agent is perceived as the type corresponding to a "Social Bliss Point" (SBP) that is represented by α. These social preferences are characterized by a social utility function $h(b_i - \alpha)$, with b_i representing an agent's perceived type. Similar to the intrinsic utility function, $h(b_i - \alpha)$ is assumed twice continuously differentiable, strictly concave, symmetric, and maximized at $b_i = \alpha$. Returning to the insider trading example, the social bliss point could be construed as the neutral asset position, with the social utility function tracking the likelihood that the player is not investigated based on how their perceived type (informed, ideal asset position) deviates from the neutral position.

The intrinsic and social utility components for a player of type t who chooses action a and is perceived as type b when the social bliss point is α are combined in a simple weighted average to form the player's total utility: [5]

$$u(a, t; b, \alpha, \lambda) = g(a - t) + \lambda h(b - \alpha) \tag{3.1}$$

Here and below, λ, the parameter that governs the degree to which an agent's utility is influenced by social considerations, is referred to as the *social preference intensity* and figures prominently in forming necessary and sufficient conditions for a conformist outcome from the model.

In Bernheim's initial development, he exogenously defined the Social Bliss

[5]The notation $f(x; y)$ represents the perspective of the bivariate function as a function of x conditional on a set of exogenous parameters y. As some of the variables move into the background, they may be dropped from the notation for convenience.

Point to be equal to unity. The chief contribution of the current work is to take a slightly different approach and, in place of exogenously defining the SBP, endogenously determine the SBP as part of the equilibrium by defining it as a function of players' actions. In particular, the present analysis adopts the notion of the SBP as corresponding to the expected action, i.e., $\alpha = E_f\left[a\left(t\right)\right]$.[6] Here, the interpretation is that individuals want to be perceived as the type who would take the average action, a characterization that can again be motivated by the insider trading example. By definition, the average action in the market is to hold each asset in proportion to its market value, which is also the optimal action for an uninformed investor (for example, from Sharpe (1964) and Lintner (1965)). As such, any investor who deviates from that position is either investing irrationally, or acting on some private information that is not publicly available. Hence, an enforcement agency such as the S.E.C. could very reasonably interpret investment positions deviant from the average action as indicative of an investor possessing some non-public information.

To close the model in the Bayesian construction, a player must form beliefs relating to the uncertainty in the game. Here, beliefs about how a player's type will be perceived are represented by an inference function, $\phi\left(b, a; \alpha, \lambda\right)$, which, conditional on social bliss point α and social preference intensity λ assign probability $\phi\left(b, a; \alpha, \lambda\right)$ to the player being perceived as type b when they take the action a. The introduction of others' actions into the social bliss point and, consequently, into each agent's utility function requires the agent to form beliefs about the distribution of the social bliss point, here represented by the measure $\pi\left(\lambda\right)$.

Given the specifications above, the individual agent's total utility maximization problem becomes:

---

[6] The current result extends to any social bliss point that is a measurable function of players' observed actions by application of a law of large numbers. For example, the social bliss point may be considered the average action actually chosen by the population of players in the game. For any finite number of players, this definition would require the distribution of beliefs $\pi(\lambda)$ to be non-degenerate. However, in the presence of a large number of players, the law of large numbers simplifies the analysis to the present case. All the basic equilibrium results go through directly, with the exception of the identification theorem (Theorem 3.3), which requires additional restrictions on the sensitivity of the social bliss point function to be noted later.

$$\max_{a \in A} E\left[u\left(a, t; \alpha, \lambda\right)\right] \tag{3.2}$$
$$= g\left(a - t\right) + \lambda \int_{\hat{\alpha} \in T} \left(\int_{b \in T} h\left(b, a; \hat{\alpha}\right) \phi\left(b, a; \hat{\alpha}, \lambda\right) db \right) d\pi\left(\hat{\alpha}\right)$$

In the current setting where the social bliss point corresponds the population expected value, the beliefs $\pi\left(\lambda\right)$ corresponds to a degenerate point distribution and the double integral is immediately reduced to a single expectation. Given this simplification, the optimization problem becomes:

$$\max_{a \in A} E\left[u\left(a, t; \hat{\alpha}, \lambda\right)\right] = g\left(a - t\right) + \lambda \int_{b \in T} h\left(b, a; \hat{\alpha}\right) \phi\left(b, a; \hat{\alpha}, \lambda\right) db \tag{3.3}$$

where $\hat{\alpha} = E_{\pi(\lambda)}\left[\alpha\right]$

3.2.1 Example: Quadratic Utility and the Spherical Case

To ground the model in a concrete functional representation, consider the simple "spherical case" example used by Bernheim to illustrate the interaction of intrinsic and social preferences as, at this stage of development, the models are little different. For the illustration, let $g\left(z\right) = -z^2$, and, $h\left(b; \alpha\right) = -\left(b - \alpha\right)^2$, then an agent's indifference curves in the (a, b) plane are defined by the equation $C = -\left(a - t\right)^2 - \lambda\left(b - \alpha\right)^2$.

As in Bernheim's example, when the social preference intensity is unity the indifference curves appear as concentric circles centered on the point (t, α), as are depicted in Figure 3.1. If the social preference intensity were to deviate from unity, these curves would appear as ellipses. Extending beyond the spherical case, the symmetry and satiation points from the current model of utility specification give rise to Bernheim's general characterization of indifference curves as:

1. horizontal at $a = t$ and symmetric around the line $a = t$.

2. vertical at $b = \alpha$ and symmetric around the line $b = \alpha$.

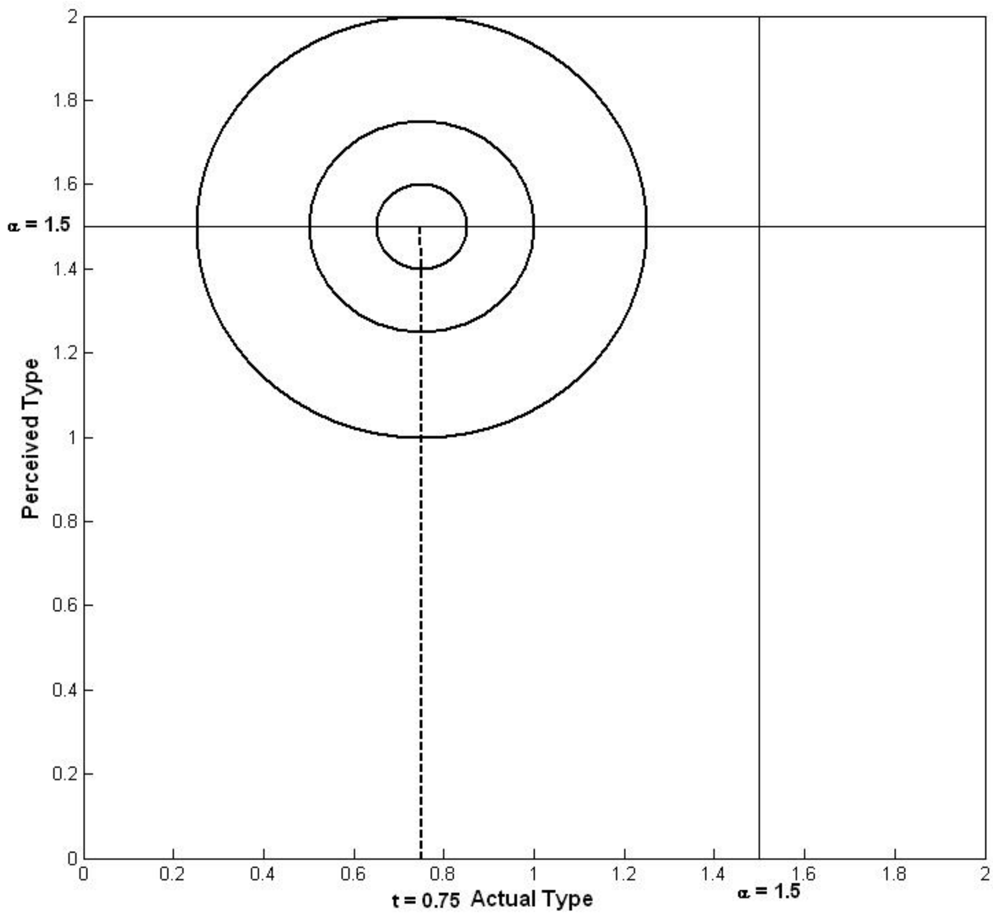

Figure 3.1: Agent's Indifference Curves

3.2.2 Non-triviality of Conformity

In Bernheim's paper, he notes that "concern over popularity does not explain conformity by itself... behavior in such a world would be observationally equivalent to that occurring in a society in which the distribution of IBPs [types] was somewhat more concentrated and in which no one cared about popularity." This observation also applies to the current setting and can be easily verified by considering the naïve inference example where:

$$\phi(b, a) = \begin{cases} 1 & \text{if } b = a \\ 0 & \text{otherwise} \end{cases} \tag{3.4}$$

The first order condition for maximizing utility under this inference function is:

$$g'\left(a^*\left(t;\alpha,\lambda\right)-t\right)+\lambda h'\left(a^*\left(t;\alpha,\lambda\right)-\alpha\right)=0 \tag{3.5}$$

Following Bernheim's lead, implicit differentiation directly implies that:

$$\frac{da^*\left(t;\alpha,\lambda\right)}{dt}=\frac{g''\left(a^*\left(t;\alpha,\lambda\right)-t\right)}{g''\left(a^*\left(t;\alpha,\lambda\right)-t\right)+\lambda h''\left(a^*\left(t;\alpha,\lambda\right)-\alpha\right)} \tag{3.6}$$

The assumed strict concavity of g and h ensures that the right hand side is strictly positive. Hence, as shown by Bernheim, direct social preferences alone do not establish conformity in the model and, further, the implied inferences are not self-sustaining as the separation of their actions ought to enable them to accurately identify each type.

3.3 Characterizing Equilibrium

Having defined preferences, beliefs, and a simple example, the next step in the analysis is to adopt an equilibrium notion and characterize the set of equilibria. Here, the natural equilibrium concept is interim Bayes Perfect Nash Equilibrium, which requires:

1. An action function, $a^*\left(t;\alpha,\lambda,\phi\right):T\rightarrow A$, such that for all $a'\in A$ and $t\in T$,

$$U\left(a^*\left(t;\alpha,\lambda,\phi\right),t;\alpha,\lambda,\phi\right)\geq U\left(a',t;\alpha,\lambda,\phi\right)$$

2. A conditional inference function, $\phi\left(b,a;\alpha,\lambda\right)$, that represents a probability distribution over the agent's inferred type, b, given their action a. The inference function is only restricted to be consistent with Bayes' Rule along the equilibrium path, though further refinements restricting off-path beliefs are discussed below in the section on equilibria with incomplete separation.

3. A set of beliefs, denoted by π that represent a probability distribution for the expected mean action, where π is asymptotically normal with mean

$E_\pi[\alpha(\lambda)] = \alpha^*(\lambda) = \int_T a^*(t; \alpha^*(\lambda), \lambda) f(t)\, dt$ and variance going to zero as I gets large.

Under Bernheim's analysis, with the exogenously fixed social bliss point, only objects satisfying conditions (1) and (2) are required conditions for a Bayes Perfect Nash Equilibrium. In adapting the theorems from that setting to the current environment, it will often be convenient to refer to an admissible action function and conditional inference function for a fixed social bliss point as an *exogenous social equilibrium.*

Having specified the model and requirements for equilibrium, the equilibrium analysis initiates by reviewing Bernheim's first main result, establishing montonicity of the optimal response function in an agent's type:

Bernheim Theorem 3.B.1 *Ceteris paribus, if $t > t'$, then the optimal response function in any equilibrium must satisfy $a^*(t; \alpha, \lambda, \phi) \geq a^*(t'; \alpha, \lambda, \phi)$.*

Proof. The proof here, which relies simply on equilibrium condition (1) coupled with the strict concavity of g, is identical to Bernheim's. The same device is used to prove Lemma 3.A.1 in the appendix and is not repeated for brevity. ■

3.4 Fully Separating Equilibrium

The analysis of conformist outcomes from the model begins by identifying the necessary and sufficient conditions for a conformist outcome to be in equilibrium. As will be discussed in greater detail in the next section, the conditions for a conformist outcome are the inverse of the necessary and sufficient conditions for a fully separating equilibrium. This section shows that incorporating an endogenous social bliss point does not change these conditions from the original setting with an exogenously defined SBP. As in Bernheim's initial development, the fully separating exogenous equilibrium for some fixed α is fully identified by an inference characterizing function $\phi_s(a)$, such that,

$$\phi(b, a) = \begin{cases} 1 & \text{if} b = \phi_s(a) \\ 0 & \text{otherwise} \end{cases} \tag{3.7}$$

By direct differentiation, the slope of the indifference curve in the (a, b) plane for an agent of type t is:

$$\frac{db}{da} = -\frac{g'(a-t)}{\lambda h'(b-\alpha)} \tag{3.8}$$

Utility optimization requires that these indifference curves be tangent to the inference characterizing function in equilibrium. I.e.,

$$\phi_s'(a;\alpha) = -\frac{g'(a-t)}{\lambda h'(\phi_s(a;\alpha)-\alpha)} \tag{3.9}$$

Further, the requirement that the inference functions are correct allows the removal of t from the equilibrium condition:

$$\phi_s'(a;\alpha) = -\frac{g'(a-\phi_s(a;\alpha))}{\lambda h'(\phi_s(a;\alpha)-\alpha)} \tag{3.10}$$

With the exception of introducing α to the notation, equation 3.10 is identical to Bernheim's equation (12) and, following his original analysis, defines a first-order differential equation for ϕ_s. Similarly, the relevant initial conditions are that the extreme types are identified exactly and, as such, have no social incentive to deviate from actions revealing their true types, i.e.,

$$\phi_s(0;\alpha) = 0, \text{ and, } \phi_s(2;\alpha) = 2 \tag{3.11}$$

In this sense, for fixed α, the analysis is no different from the Bernheim setting in that equation 3.10 with initial conditions 3.11 defining a set of two first-order differential equations in the (a, b) plane, the first starting from $(0, 0)$ and moving northeast and the second starting from the point $(2, 2)$ and moving southwest, that must meet at a unique point. The results from Bernheim's analysis for any fixed SBP follow directly, with only minor adjustments to accommodate the non-centrality of the SBP.

Following Bernheim, one more piece of technicalia must be developed to define the equilibrium. For agents with this type $t \leq \alpha$, consider the range of the inference characterizing function by defining $\underline{A} = \phi_s^{-1}([0, \alpha])$ to represent the set of all choices made by individuals with $t \in [0, \alpha]$. By monotonicity, $\underline{A}$ is an interval $[0, \underline{a}]$ over which the existence and uniqueness of ϕ_s is established by standard arguments. Further, over $\underline{A}$, $\phi_s(a) \leq a$, yielding the key implication that $\underline{a} \geq \alpha$. The analysis continues by considering the other portion of the type space, i.e., the agents with types $t \geq \alpha$. Correspondingly, define the interval $[\bar{a}, 2] = \bar{X} = \phi_s^{-1}([\alpha, 2])$ and parallel arguments yield the implication that $\bar{a} \leq \alpha$.

The apparatus in the previous paragraph provides a framework for establishing simple conditions for the inference characterizing function to be well defined. More precisely, a fully separating equilibrium requires that a unique type be assigned to each action and, based on the previous paragraph, is equivalent to the condition that:

$$\underline{a} = \bar{a} = \alpha \tag{3.12}$$

The main result of this section is to establish that the model admits a fully separating equilibrium if and only if the social preference intensity is lower than some critical level. Having established the characterization of a fully separating exogenous social equilibrium, the argument establishing necessary and sufficient conditions for such an outcome begins by identifying a unique α^* consistent with a fully separating equilibrium (this result corresponds to Lemma 3.1, below). After identifying a unique fully separating endogenous social equilibrium, the next step of the argument is to establish that, for each fixed α, there is a critical level $\lambda^*(\alpha)$ such that, for any $\lambda \leq \lambda^*(\alpha)$, there exists a fully separating equilibrium and, for any $\lambda > \lambda^*(\alpha)$ there is no fully separating equilibrium (which is proved below in theorem 3.B.2). The argument closes by showing that, even though these critical levels are identified for each fixed α, they are constant in that parameter, so that $\lambda^*(\alpha) = \lambda^*(\alpha')$ (this is established as Theorem 3.1).

Lemma 3.1 *Given the optimal response function $a^*(t; \alpha, \lambda)$ with a consistent inference function $\phi(b, a; \alpha, \lambda)$ for every $\alpha \in T$, there exists a unique α^* such that $\alpha^* = E_f[a^*(t; \alpha^*, \lambda, \phi)]$.*

Lemma 3.1 establishes that, over all the exogenous social equilibria, there is a unique endogenous social equilibrium with a social bliss point satisfying equilibrium condition (3). Lemma 3.1 is a direct result of the contraction mapping theorem, with the crux of the proof showing that the mapping $\gamma_\lambda(\alpha) : \alpha \mapsto \int_T a^*(t;\alpha,\lambda) f(t)\, dt$ is a contraction mapping for any λ. Intuitively, the argument states that the degree to which the average player's action changes in response to a change in the SBP is strictly less than the change in the SBP. The argument obtaining a less than or equal to relationship is derived through the additional dependence of utility on the IBP, while the argument for a strictly less than relationship exploits the fixed end points of the differential equation defining ϕ_s. The mathematical details are rather tedious and so are deferred to the appendix.

The second phase of the argument assumes the existence of an equilibrium social bliss point and establishes necessary and sufficient conditions for a separating equilibrium given that SBP. This claim is verified by first establishing that there exists some level of social preference intensity (λ) where the inference characterizing function satisfies the separating equilibrium conditions 3.10, 3.11, and 3.12. Next, it is shown that there is some level of social preference intensity where the inference characterizing function fails to satisfy these conditions. As in Bernheim's analysis, the general result is shown to be monotonic in λ for fixed social bliss points and can be stated exactly as he does in his Theorem 2.

Bernheim Theorem 3.B.2 *For each α, there exists $\lambda^*(\alpha) > 0$ such that a fully separating equilibrium exists if and only if $\lambda \leq \lambda^*(\alpha)$.*

The proof in the appendix follows Bernheim's reasoning closely, though it is somewhat more difficult to establish the monotonicity result he uses to complete the proof. The results are finalized by dealing with the fact that the SBP is not insensitive to the social preference intensity in Theorem 3.1 but still identifies a unique critical value for the social preference intensity to obtain a conformist outcome.

Theorem 3.1 *There exists a unique $\lambda^* > 0$ such that a fully separating equilibrium exists if and only if $\lambda \leq \lambda^*$*

This result is again illustrated using the spherical example developed in

Bernheim. In the spherical example, equation 3.10 becomes:

$$\phi_s'(a;\alpha) = -\left(\frac{1}{\lambda}\right)\left[\frac{a-\phi_s(a;\alpha)}{\alpha-\phi_s(a;\alpha)}\right] \tag{3.13}$$

To find a critical value that ensures the existence of a fully separating equilibrium, analyze equation (6) as a linear dynamical system in (t, x):

$$\begin{bmatrix} {dt}/{d\tau} \\ {dx}/{d\tau} \end{bmatrix} = \begin{bmatrix} x-t \\ \lambda(\alpha-t) \end{bmatrix} = \begin{bmatrix} -1 & 1 \\ -\lambda & 0 \end{bmatrix}\begin{bmatrix} t-\alpha \\ x-\alpha \end{bmatrix} = A\begin{bmatrix} t-\alpha \\ x-\alpha \end{bmatrix} \tag{3.14}$$

In equation 3.14, τ is an indexing variable and the existence a fully separating equilibrium is equivalent to the matrix $\mathbf{A}$ having real eigenvalues, a condition that is entirely independent of the social bliss point (note: this result is unique to the spherical case). As the matrix $\mathbf{A}$ in this setting is identical to Bernheim's, the critical value λ^* for satisfying existence of a fully separating equilibrium is easily seen to be $\lambda^* = {1}/{4}$ for all social bliss points.

Figure 3.2, below, plots the two ODE's starting from the extreme points for the spherical case where $\lambda \in \{0, {1}/{10}, {1}/{4}, 1\}$ when the social bliss point is 1.5. It is interesting to note that, although perfect symmetry is lost, the inference characterizing function in the northeast segment of the graph is a scaled version of that in the southwest. The symmetry in this example arises from the homothetic indifference curves.

3.5 Equilibria with Incomplete Separation

As is common in rich signaling games, in the current setting there is a great multiplicity of equilibria with incomplete separation due to a lack of restrictions on off equilibrium path beliefs in the Bayesian equilibrium. As a pre-emptive step in addressing this multiplicity and ruling out equilibria characterized by implausible beliefs, this section adopts the D1 criterion as an off equilibrium path belief restriction. With this equilibrium refinement in hand, the analysis proceeds to characterize the pooling equilibria as having exactly the same properties estab-

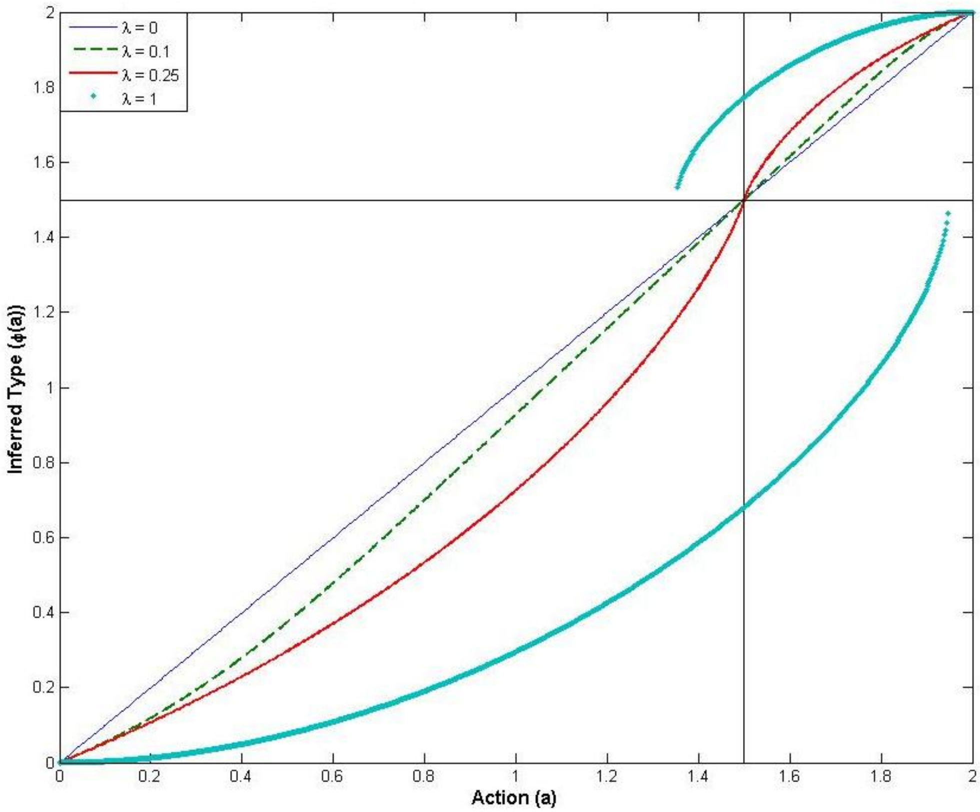

Figure 3.2: Separating Inference Function in Spherical Case when $\alpha = 1.5$

lished in Bernheim's initial development. The first significant contribution in this section in extending Bernheim's initial development shows that generalizing the model to incorporate an endogenous social bliss point does not introduce further multiplicity in the set of equilibria. A further result is established showing that the endogenous social bliss point and the multiplicity of equilibria in Bernheim's initial development are coupled one-for-one, so that each equilibrium is uniquely identified by its social bliss point. As such, by identifying an exogenous social bliss point (for example, one that maximizes a social welfare function), one is able to identify a unique equilibrium with incomplete separation.

In order to reasonably constrain the set of equilibria, Bernheim adopts the D1 criterion for restricting off-equilibrium path beliefs. An intuitive justification for invoking this refinement is based on a forward induction argument, where the D1 criteria is interpreted as an extension of the Intuitive Criteria from Cho and

Kreps (1987) in a manner similar to that of Divinity Banks and Sobel (1987). A theoretical justification is that the D1 criterion ensures the equilibrium is strategically stable, in the sense of Kohlberg and Mertens (1986). Further, the D1 criterion ensures conformity does not arise primarily due to unreasonable off-path beliefs. In particular, the refinement rejects pooling equilibria whenever a fully separating equilibrium is attainable and has been shown to preclude interior pooling in cases where indifference curves satisfy the single crossing property (which is not the case here given the dual symmetries of the indifference curves), as established by Cho and Sobel (1990).

3.5.1 The Bernheim Results

The analysis begins by introducing some notation to review the Bernheim results. Define: $T(a) = \{t \in T | a^*(t) = a\}$, and, $t_l(a) = \inf T(a)$, and, $t_h(a) = \sup T(a)$

Note that the monotonicity established by Bernheim's Theorem 3.B.1 implies that $T(a)$ is an interval and, since each individual type has measure zero, can be written as a closed set including its endpoints: $T(a) = [t_l(a), t_h(a)]$. It will also be helpful to recall the definition of an exogenous social equilibrium as any action and inference function satisfying equilibrium conditions (1) and (2) for some exogenously fixed α.

Bernheim's theorems related to partial separation generalize directly to the current context as none of the arguments in his proofs rely directly on the centrality of the social bliss point.

Bernheim Theorem 3.B.3 *For fixed α, if $\lambda > \lambda^*$, then for any exogenous social equilibrium that satisfies the D1 criterion, there exists at most one $a_p \in A$ such that $t_l(a_p) < t_h(a_p)$, and it satisfies $\alpha \in T(a_p)$.*

Bernheim Theorem 3.B.4 *For fixed α and any given $a_p \in A$, there is at most one central pooling exogenous social equilibrium (a_p, t_l, t_h).*

These two theorems characterize pooling equilibria as consisting of a single pooling point and, for each pooling point, a fixed set of types participating in the pool. Outside of the pool, i.e., for agents with types $t \notin T(a_p)$ behavior is characterized by the same inference characterizing function discussed in the fully separating equilibrium, that is to say, $a^*(t,\alpha) = \phi_s^{-1}(t;\alpha)$.

3.5.2 Implications of Endogenous Social Bliss Point

While the Bernheim results establish candidate exogenous social equilibria, they do not address equilibrium condition (3) in the extended model. Here, condition (3) can be exploited to further restrict the set of equilibria and doing so yields the first significant result of this section:

Theorem 3.2 *If $\lambda > \lambda^*$, then conditional on the population average strategy, the unique social equilibrium with incomplete separation satisfying the D1 criterion is characterized by a single central pool at $a_p^* = \alpha^* + \varepsilon(\alpha^*,\lambda)$, where:*

$$\varepsilon(\alpha^*,\lambda) = \tag{3.15}$$
$$\left(\frac{\int_0^{t_l(a_p)} \left(\alpha^* - \phi_s^{-1}(t;\alpha^*,\lambda)\right) dF(t) + \int_{t_h(a_p)}^{2} \left(\alpha^* - \phi_s^{-1}(t;\alpha^*,\lambda)\right) dF(t)}{P\left(t \in \left[t_l(a_p), t_h(a_p)\right]\right)}\right)$$

Theorem 3.2 identifies the pooling point as a function of the population average action and an additive perturbation. This perturbation is shown in Theorem 3.3 to be continuous in the pooling point, an implication that is extended by the Intermediate Value Theorem argument to show that unique equilibrium can be identified where $\varepsilon(\alpha^*,\lambda) = 0$, which is perhaps the most surprising result in the current work.

Theorem 3.3 *If $\lambda > \lambda^*$, then there exists a unique social equilibrium with incomplete separation satisfying the D1 criterion where the single central pool is located at the population average strategy, i.e., where $\varepsilon(\alpha^*,\lambda) = 0$.*

Note that much of the variability in the model arises from the broad set of allowable distributions and this can muddle the result. Additional insight into the restrictions imposed on pooling equilibria by the endogenized social bliss point is provided by the next corollary, which directly implies a similar result would hold in the initial development if Bernheim were to assume the pooling action satisfies some social welfare function.

Corollary 3.1 *If the distribution over types is symmetric around 1, then the unique social equilibrium consistent with pooling on the expected action is* $a_p^* = \alpha^* = 1$.

3.6 The Social Preference Intensity

To this point, the social preference intensity parameter, λ, has been taken as exogenously defined and constant across types. While there is no natural mechanism for endogenizing this parameter, this section considers the equilibrium effects of a generalized social preference intensity that can depend upon a player's type. To initiate the analysis, consider the general statement of the agent's utility maximization problem as:

$$\max_{a \in A} E\left[u\left(a, t; \hat{\alpha}, \lambda\left(t\right)\right)\right] = g\left(a - t\right) + \lambda\left(t\right) \int_{b \in T} h\left(b, a; \hat{\alpha}\right) \phi\left(b, a; \hat{\alpha}, \lambda\left(t\right)\right) db \tag{3.16}$$

where $\alpha = E_f\left[a^*\left(t; \alpha, \lambda\left(t\right)\right)\right]$ and now $\lambda : T \rightarrow \Lambda \subset \mathbb{R}_+$ is a function assigning a social preference intensity to each individual's type.

An exhaustive analysis of models of this type is beyond the scope of the current work, so this section focuses on presenting comparative analysis of several particular specifications for the social preference intensity function. For tractability in the current analysis, it will be easiest to think of the SBP as a constant here, with generalizations readily available by invoking the contraction property.

3.6.1 Continuous, Monotonic Social Preference Intensities

To further narrow the problem under consideration, the analysis begins by considering the case where the social preference intensity function is twice continuously differentiable, symmetric, and monotonically decreasing in $|t - \alpha|$. Here, the continuity and symmetry properties are intended to mimic the other properties of the utility function while the monotonicity assumption is intended to capture a stylized fact that individuals far from the social norm are less affected by social concerns.

The results in this case are no different from the above analysis. If the social preference intensity exceeds a threshold level for some player types, then a single pooling equilibrium will be established, with the pooling action identified uniquely by the social bliss point. If the social preference intensity is bounded below this threshold level, a fully separating equilibrium obtains.

3.6.2 Discontinuous Functions of Types

Perhaps a more interesting specification of the social preference intensity is to consider a setting where the social preference intensities take the form of a step function that is monotonically *increasing* in $|t - \alpha|$. In this setting, agents who are far from the social norm actually lend greater credence to social considerations in their preferences than those who are near the norm. For simplicity, suppose there are only two discretely different social preference intensities, λ_l for agents with types in a neighborhood of the social bliss point denoted Δ and λ_h for agents with types far from the social bliss point.

Such a setting will, in general, directly yield one or two centralized conformist pools, regardless of whether the social preference intensities are lower than the threshold value. The pooling result arises because the shift in social preference intensities leads to a discontinuous jump in the ODE defined by equation (3). This discontinuity takes the form of a kink in the inference characterizing function, causing the function to turn sharply towards the social bliss point. Further, this kink in inferences can be exploited by agents whose types are on the border of Δ to mimic agents with lower social preference intensities. These pools of perceived apathy will separately form on both sides of the social bliss point unless the social preference intensity is so high, or the neighborhood Δ is so small, that the pools happen to spill into each other.

In contrast, suppose the social intensity function were a step function that is monotonically decreasing in $|t-\alpha|$ as in the previous section. In this case, the kink in the inference characterizing function would actually shift the curve further away from the 45^o line, a kink that would not provide the opportunity for agents to mimic another types action unless the social preference intensity exceeds some threshold value.

3.7 Potential Extensions and Applications

Given the theoretical results above, there are a number of potential avenues for developing applications of these findings. Here, I briefly mention some directions these developments might take.

3.7.1 Uncertain Bliss Points

Given an exogenously specified social preference intensity, a plausible dynamic for endogenizing social preference intensities could be defined by generalizing the intrinsic bliss point so that the individual's type and their IBP satisfies the monotone likelihood ratio property. For example, instead of a player's intrinsic bliss point being identically equal to their type, their bliss point might be randomly distributed but centered at their type with fixed variance σ. Here, that variance represents the degree to which an individual is certain of their own intrinsic bliss point.[7]

An interesting result obtains in comparative statics when the variance of the intrinsic bliss point is allowed to vary. In particular, as the variance increases, Jensen's inequality demands that intrinsic utility drop simply due to greater variability in outcomes. However, this drop in intrinsic utility has an effect equivalent to that of an increase in the social preference intensity, yielding the result that, as uncertainty over individual preferences increases, the individual will be more affected by fads and fashions. The interaction is fairly complicated, but it could provide a nice mechanism for motivating comparative statics analysis relating to

[7]Many thanks to Jacob LaRiviere for recommending this approach to motivating the current analysis.

the behavior of insiders relative to uninformed agents.

3.7.2 Alternate Social Statistics as Modal Behavior

While the development of this piece centers on the social bliss point as the expected action, it may be interesting to look at other social statistics such as the median (particularly relevant to political contexts) or mode. The current results can be immediately extended whenever an argument can be made that the believed distribution of the social bliss point becomes degenerate with a large number of players. This condition is not innocuous in application to social bliss points identified as sample social statistics, as it effectively precludes ex post equilibria in settings with a finite number of players.

While Bernheim's model can be easily extended to allow for multi-modal pooling behavior by introducing several local maxima to the social utility function, such an extension in the setting with an endogenized social bliss point is not particularly obvious. One approach might be to consider a local interaction type of setting, where a player only really cares about those players with types perceived to be near their own. In this model, the social bliss point would become a function of the individual's type. While separation results are likely similar to the above, pooling behavior in such a model would be chaotic, with many pools forming at arbitrary locations that depend on the locations and sizes of other pools.

3.7.3 Signaling Social Preference Intensities

Another interesting extension would be to incorporate an individual's social preference parameter into the player's type and be imperfectly observed by those around you. For example, the player's social preference intensity could be distributed as a martingale with various measurability restrictions. In such a setting, a player's private information consists of both their intrinsic bliss point and their social preference intensity. Despite two dimensions of private information, each player-type will choose a single action to maximize their utility. A significant challenge lies in consistently inferring their two-dimensional type from their single action, which would create a very interesting interaction requiring further equilibrium refinements.

3.7.4 Applications

As for applications, a significant motivation in developing the endogenized social bliss point was to render the model more accessible to application and the marginal contribution facilitating equilibrium identification allows for an estimable model.

One potential application to industrial organization lies in modeling signals related to product characteristics. Suppose a firm observes the quality of their product perfectly and sends a signal to the public about their product's quality. It's costly to mis-represent their product quality, in the form of returned product and customer complaints. However, the firm's sales benefit from making their product appear closer to the average product on the market. An empirical application might look at different industries with varying levels of interactions in the market to see if coordination is more likely to arise in settings where there is a low cost of mis-representing product quality or a high cost for a non-standard product.

Another interesting application would be to adapt the current model to a signaling problem faced by forecasters. In this game, the forecasters announce public forecasts to signal their private information to a population of customers. The forecasters gain utility from announcing a forecast consistent with their private information (represented by their IBP) but the customers, who lack any private information relating to the true state of the world, award private contracts based on the forecaster's distance from an aggregated forecast (the SBP). The social preference intensity would be determined by the value of the available private contracts and the degree of uncertainty regarding the true state of the world. Given the relationship between the social preference intensity and the degree to which forecasters would tilt their announced forecasts toward the mean, these features could organize a number of results observed in analyzing the term structure of the cross-sectional variation in forecasts.

3.8 Conclusion

The model developed above provides a reasonable context for exploring conformist behavior in a variety of settings where individuals possess varied tastes but also consider social consequences when acting with incomplete information. The current model is particularly adapted to characterizing social considerations can that guide an individual's behavior to move towards the cultural mean action but can be applied in any setting where the individual's social considerations are centered at a location identified by other agent's actions. Further, the model illustrates that, in a setting where social considerations dominate, individuals will suppress their individual preferences and behave in a completely conformist manner.

Given the earlier results from Bernheim (1994), the conformist outcome is non-surprising as the basic results from his exercise are expected to obtain. However, it is surprising that, despite introducing a generalization to the original model, no new multiplicity is introduced to the set of equilibria and the necessary and sufficient conditions for a conformist outcome are unchanged. Further, the connection between the multiplicity of equilibria from the initial Bernheim development and the indeterminant endognenous social bliss point is entirely unexpected. Through this mechanism, the current setting allows for a sharper characterization of how conformity is expressed and finds a unique equilibrium where conformity is actually consistent with the social norm. While non-norm based pooling equilibria survive the extension, the set of equilibria is no larger than in Bernheim's original model and additional insight is gained into what allows these apparently inconsistent equilibria to exist.

Appendix: Proofs

Proof of Lemma 3.1

Given the optimal response function $a^* (t; \alpha, \lambda)$ with a consistent inference function $\phi (b, a; \alpha, \lambda)$ for every $\alpha \in T$, there exists a unique Social Bliss Point, α^*, such that $\alpha^* = E_f [a^* (t; \alpha^*, \lambda, \phi)]$.

Proof. As mentioned in the main section, the crux of the argument is showing that $\gamma_\lambda (\alpha) : \alpha \mapsto \int a^* (t; \alpha, \lambda) f (t) \, dt$ is a contraction mapping. Intuitively, the argument is that the degree to which a player's strategy is expected to change in response to a change in the expected mean action of the other players is strictly less than the shift in the mean action. Mathematically, it requires verifying that for all α' and α'' in T:

$$\left| \int a^* (t; \alpha', \lambda) f (t) \, dt - \int a^* (t; \alpha'', \lambda) f (t) \, dt \right| < |\alpha' - \alpha''| \tag{3.A.1}$$

Lemma 3.A.1 *For all $t \in T$ and all $\alpha' \leq \alpha''$, $a^* (t; \alpha', \lambda) \leq a^* (t; \alpha'', \lambda)$.*

Proof. The argument here is similar to the Bernheim proof of his Theorem 1, but works primarily on the social utility function h rather than the intrinsic utility function g. Let r be the intrinsic utility associated with choosing $a = a^* (t; \alpha)$ and let r' be the level of intrinsic utility associated with $a' = a^* (t; \alpha')$. Assume $a' > a$ in an equilibrium, which requires:

$r + h (a - \alpha) \geq r' + h (a' - \alpha)$, and, $r' + h (a' - \alpha') \geq r + h (a - \alpha')$

Adding these two inequalities gives:

$$h (a' - \alpha') - h (a - \alpha') \geq h (a' - \alpha) - h (a - \alpha) \tag{3.A.2}$$

Now using the Bernheim trick of applying the Fundamental Theorem of Calculus twice and using the strict concavity of h:

$$\begin{aligned}&[h(a'-\alpha')-h(a-\alpha')]-[h(a'-\alpha)-h(a-\alpha)]\\&=\int_a^{a'}h'(w-\alpha')-h'(w-\alpha)\,dw=\int_a^{a'}\int_\alpha^{\alpha'}h''(w-v)\,dwdv<0\end{aligned}\tag{3.A.3}$$

However, this last inequality contradicts 3.A.2 and Lemma 3.A.1 is proved.

■

Lemma 3.A.3 *For any $t\in T$ and $\alpha'\leq\alpha''$, let $a'=a^*(t;\alpha',\lambda)$ and $a''=a^*(t;\alpha'',\lambda)$, then $a''-a'\leq\alpha''-\alpha'$.*

Proof. This lemma holds simply by evaluating the "income" and "substitution" effects associated with the shift in the social bliss point. First, suppose by contradiction that $a''-a'>\alpha''-\alpha'$. From an agent's perspective in (a,b) space, the shift in SBP effectively corresponds to a translation of their utility functions and can be analyzed equivalently to a shift in the budget set (here the inference characterizing function). Here, then, reacting exclusively to the shift in the SBP corresponds to the income effect and can be compensated entirely by exactly translating the inference characterizing function by the same magnitude as the indifference curves. However, such a translation gives rise to substitution effects, and agents will substitute intrinsic utility for social utility. Hence, the only way an agent can overcompensate for a shift in the SBP is if the substitution effect were somehow negative, contradicting the concavity assumptions of both the intrinsic and social utility functions. ■

Lemma 3.A.3 *There is some $\varepsilon>0$ so that:*

$$\left|\int a^*(t;\alpha',\lambda)\,f(t)\,dt-\int a^*(t;\alpha'',\lambda)\,f(t)\,dt\right|<|\alpha'-\alpha''|-\varepsilon\tag{3.A.4}$$

Proof. The proof here revolves around the fixed end points, and proceeds by computing the integrals on the left hand side over the range $[0,\delta)$ and $(2-\delta,2]$. Because $\phi_s(0)$ and$\phi_s(2)=1$, the integral will be strictly less than the value required to make $\varepsilon=0$:

$$|\alpha'-\alpha''|\,P\left(t\in[0,\delta)\cup(2-\delta,2]\right)\tag{3.A.5}$$

The result then follows by Jensen's Inequality.

■

■

Proof of Bernheim Theorem 3.B.2

For each α, there exists $\lambda^*(\alpha) > 0$ such that a fully separating equilibrium exists if and only if $\lambda \leq \lambda^*(\alpha)$.

Proof. These arguments follow Bernheim's proof almost exactly. The only adaptation to the current setting, while admittedly non-trivial, is to re-center each of the steps at α in place of 1 and adding a few minimum and maximum operations. These devices are here developed explicitly primarily for illustrative purposes and this detail is not included in the below proofs of other theorems from the Bernheim's original work. For the purposes of these devices, assume $\alpha > 1$ since the case of $\alpha = 1$ is established by Bernheim and the case $\alpha < 1$ will follow by symmetry.

Step 1: For λ sufficiently large, no fully separating equilibrium exists.

To verify this, choose $\lambda > \max\left\{\frac{g(0)-g(\alpha)}{h(\alpha)-h(0)}, \frac{g(2)-g(\alpha)}{h(\alpha)-h(2)}\right\}$. If $a^*(\alpha;\alpha,\lambda) \leq \alpha$, type 0 agents would have an incentive to imitate type α agents. However, if $a^*(\alpha;\alpha,\lambda) > \alpha$, then type 2 agents would have an incentive to imitate type α agents. Since pooling would arise for one of these extreme types, no separating equilibrium exists.

Step 2: If $\lambda > 0$ is sufficiently small, a fully separating equilibrium does exist, i.e., for small λ, $\bar{a} = \underline{a} = \alpha$.

Step 2a: There exists a small enough λ so that players with types below the SBP play fully separating strategies.

Choose some $\theta > \alpha$ and define the line segment $B(a) = (\alpha - \theta) + \frac{\theta}{\alpha}a$ over the interval $\left[\frac{(\theta-\alpha)}{\theta}\alpha, \alpha\right]$ so that $B\left(\frac{(\theta-\alpha)}{\theta}\alpha\right) = 0$ and $B(\alpha) = \alpha$. Now, there is some $K > 0$ such that, for $a \in [(\theta-\alpha)\, a/\theta, \alpha)$, $G(a) \equiv -\frac{g'(a-B(a))}{h'(B(a))} > K$.[8] Since $a > B(a)$ and $B(a) < \alpha$, $G(a)$ is strictly positive for $a < \alpha$,. The claim can only be false, then, if there is some sequence $\langle a_k \rangle_{k=0}^{\infty}$ such that $\lim_{k\to\infty} G(a_k) = 0$. Without loss of generality, suppose the sequence converges to a single limit point $\hat{a}$. Suppose $\hat{a} < \alpha$, then since G is continuous, $G(\hat{a}) = 0$ contradicts the strict positivity of $G(a)$ for $a \in [(\theta-\alpha)\,\alpha/\theta, \alpha)$. Now suppose $\hat{a} = \alpha$, then the limit of $G(a_k)$ can be computed using L'Hospital's rule: $\lim\limits_{a\to\alpha} G(a) = -\frac{(1-\theta)g''(0)}{\theta h''(\alpha)} > 0$, this contradicts the hypothesis that the limit of $\langle G(a_k) \rangle_{k=0}^{\infty}$ equals 0. These two contradictions establish the claim.

Step 2b: There exists a small enough λ so that players with types above the SBP play fully separating strategies.

The argument that $\underline{a} = \alpha$ follows by a similar device but is more confusing because it all operates in reverse. Choose some $\underline{\theta} > \max\left\{2-\alpha, \frac{\alpha^2}{2-\alpha}\right\}$and define the line segment $B(a) = \alpha - \frac{\underline{\theta}(2-\alpha)}{\alpha} + \frac{\underline{\theta}(2-\alpha)}{\alpha^2} a$ over the interval $\left[\alpha, \frac{(\theta+\alpha)}{\underline{\theta}}\alpha\right]$ so that $B\left(\frac{(\underline{\theta}+\alpha)}{\underline{\theta}}\alpha\right) = 2$ and $B(\alpha) = \alpha$. Now, there is some $\underline{K} > 0$ such that, for $a \in (\alpha, (\underline{\theta}+\alpha)\,\alpha/\underline{\theta}]$, $G(a) \equiv -\frac{g'(a-B(a))}{h'(B(a))} < \underline{K}$, or equivalently, that there exists a $\underline{k} > 0$ so that $H(a) \equiv -\frac{h'(B(a))}{g'(a-B(a))} > \underline{k}$. Since $a < B(a)$ and $B(a) > \alpha$ for $a > \alpha$,$H(a)$ is strictly positive. The claim can only be false, then, if there is some sequence $\langle a_k \rangle_{k=0}^{\infty}$ such that $\lim_{k\to\infty} H(a_k) = 0$. Without loss of generality, suppose the sequence converges to a single limit point $\hat{a}$. Suppose $\hat{a} > \alpha$, then since H is continuous, $H(\hat{a}) = 0$ contradicts the strict positivity of $H(a)$ for $a \in (\alpha, (\underline{\theta}+\alpha)\,\alpha/\underline{\theta}]$. Now suppose $\hat{a} = \alpha$, then, as above, the limit of $H(a_k)$ can be computed using L'Hospital's rule: $\lim\limits_{a\to\alpha} H(a) = -\frac{\underline{\theta}(2-\alpha)(\alpha)}{(\alpha^2-\underline{\theta}(2-\alpha))g''(0)} > 0$, this contradicts the hypothesis that the limit of $\langle G(a_k) \rangle_{k=0}^{\infty}$ equals 0. These two contradictions establish the claim.

Step 2c: The hypothesis that there is a non-zero λ so that $\bar{a} = \alpha = \underline{a}$ is established by choosing λ so that $\lambda \max\left\{\frac{\theta}{\alpha}, \frac{\alpha^2}{\underline{\theta}(2-\alpha)}\right\} < \min\{\underline{K}, \underline{k}\}$ and showing that:

[8] It may be worth noting here that the initial proof published by Bernheim in the JPE was victim of a typographical error in his equation (B5) that render his equations (B6) and (B7) to be incorrectly signed. His equation (B5) should have been defined as here, which provides the shortest route to a corrected proof.

(i) $\phi_s(a) > B(a)$ for all $a \in \left[\frac{(\theta-\alpha)\alpha}{\theta}, \alpha\right)$, and,

(ii) $\phi_s(a) < B(a)$ for all $a \in \left(\alpha, \frac{(\theta+\alpha)\alpha}{\theta}\right]$.

To prove (i), note that since $\frac{(\theta-\alpha)\alpha}{\theta} > 0$, $\phi_s\left(\frac{(\theta-\alpha)\alpha}{\theta}\right) > B\left(\frac{(\theta-\alpha)\alpha}{\theta}\right) = 0$. Now, suppose there exists some $a' \in \left(\frac{(\theta-\alpha)\alpha}{\theta}, \alpha\right)$ such that $\phi_s(a') < B(a')$, then there exists some $a'' \in \left(\frac{(\theta-\alpha)\alpha}{\theta}, a'\right)$ such that $\phi_s(a'') = B(a'')$ and $\phi_s{}'(a'') \leq B'(a'')$. However, this contradicts the result from step 2a that:

$$\phi_s{}'(a) = G(a)/\lambda > K/(\alpha\lambda) > \theta = B'(a)$$

Hence, $\phi_s(a)$ must remain above $B(a)$ when $a < \alpha$, implying that $\alpha \leq \bar{a} \leq \alpha$.

Claim (ii) is proven by first observing that since $\frac{(\theta+\alpha)\alpha}{\theta} < 2$, $\phi_s\left(\frac{(\theta+\alpha)\alpha}{\theta}\right) < B\left(\frac{(\theta+\alpha)\alpha}{\theta}\right) = 2$. Now, suppose there exists some $a' \in \left(\alpha, \frac{(\theta+\alpha)\alpha}{\theta}\right)$ such that $\phi_s(a') > B(a')$, then there exists some $a'' \in (\alpha, a')$ such that $\phi_s(a'') = B(a'')$ and $\phi_s{}'(a'') \leq B'(a'')$. However, this derivative property contradicts the result from step 2b that:

$$\phi_s{}'(a) < \lambda/(\underline{k}) < \frac{\theta(2-\alpha)}{\alpha^2} = B'(a)$$

Hence, $\phi_s(a)$ must remain below $B(a)$ when $a > \alpha$, implying that $\alpha \leq \underline{a} \leq \alpha$. This argument completes the proof of Step 2.

Step 3: The present goal is to prove a monotonicity result to the effect that, if $\lambda' \leq \lambda''$ and there is a fully separating equilibrium for $\lambda = \lambda''$, then there exists a fully separating equilibrium for $\lambda = \lambda'$.

The argument proceeds by establishing two claims.

Claim 1: If $\lambda' \leq \lambda''$, then $\int_0^2 (\phi_s(a;\alpha,\lambda') - a)^2 da \leq \int_0^2 (\phi_s(a;\alpha,\lambda'') - a)^2 da.$

Claim 2: If $\lambda' \leq \lambda''$, and $\phi_s(\alpha^*(\lambda'')) = \alpha^*(\lambda'')$, then $\phi_s(\alpha^*(\lambda')) = \alpha^*(\lambda')$

Claim 1 is a direct consequence of the first order condition 3.10. As λ becomes small, the derivative of the inference characterizing function ϕ_s increases over its entire domain. Since the inference characterizing function is bounded to be less than or equal to the identity function, the distance between ϕ_s and the identity function decreases globally.

Claim 2 also provides the key to establishing Theorem 3.1, below. It also holds as a consequence of first order equilibrium condition 3.10 and the contraction property established by Lemma 3.1. The first order equilibrium condition ensures the ODE will attain a sufficiently high slope to reach the SBP (since $h'(0) = 0$), and the contraction property ensures the existence of a unique target SBP to which the ODE converges.

■

Proof of Theorem 3.1

There exists a unique $\lambda^* > 0$ such that a fully separating equilibrium exists if and only if $\lambda \leq \lambda^*$.

Proof.

The result here is fairly direct from the previous theorem and Lemma 3.1 and the result could be considered more of a corollary. As shown in Theorem 3.B.2, there exists a $\lambda^*(\alpha) > 0$ for any α such that a separating equilibrium exists if and only if $\lambda \leq \lambda^*(\alpha)$. Connecting this theorem with Lemma 1, that under a complete specification there is exactly one equilibrium social bliss point, α^*, an immediate result is that if $\lambda^* = \lambda^*(\alpha^*)$, a fully separating equilibrium obtains if and only if $\lambda \leq \lambda^*$.

■

Proof of Theorems 3.B.3 and 3.B.4

Theorem 3.B.3: For fixed α, if $\lambda > \lambda^*$, then for any exogenous social equilibrium that satisfies the D1 criterion, there exists at most one $a_p \in A$ such that $t_l(a_p) < t_h(a_p)$, and it satisfies $\alpha \in T(a_p)$.

Theorem 3.B.4: For fixed α and any given $a_p \in A$, there is at most one central pooling exogenous social equilibrium (a_p, t_l, t_h).

Proof. Since these theorems are stated for fixed α, the proofs from Bernheim are directly applicable. An extraordinary amount of tedium would be needed to identify and remedy all issues like those addressed in Theorem 3.B.2, but nowhere in his analysis is the centrality of the SBP required.

■

Proof of Theorem 3.2

If $\lambda > \lambda^*$, then conditional on the population average strategy, the unique social equilibrium with incomplete separation satisfying the D1 criterion is characterized by a single central pool at $a_p^* = \alpha^* + \varepsilon(\alpha^*, \lambda)$, where:

$$\varepsilon(\alpha^*, \lambda) = \left(\frac{\int_0^{t_l(a_p)} \left(\alpha^* - \phi_s^{-1}(t; \alpha^*, \lambda)\right) d\pi(t) + \int_{t_h(a_p)}^2 \left(\alpha^* - \phi_s^{-1}(t; \alpha^*, \lambda)\right) d\pi(t)}{P\left(t \in [t_l(a_p), t_h(a_p)]\right)} \right)$$

Proof. The result follows immediately by applying equilibrium condition (C) to the intersection of the sets of equilibria established in Bernheim's Theorems (3) & (4). Writing condition (C) in integral form yields:

$$\begin{aligned}
\alpha^* =& E_\pi\left[a^*(t; \alpha^*, \lambda)\right] = \int_T a^*(t; \alpha^*, \lambda)\, d\pi(t) \\
=& \int_0^{t_l} a^*(t; \alpha^*, \lambda)\, d\pi(t) + \int_{t_l}^{t_h} a^*(t; \alpha^*, \lambda)\, d\pi(t) + \int_{t_h}^2 a^*(t; \alpha^*, \lambda)\, d\pi(t) \\
=& \int_0^{t_l} \phi_s^{-1}(t; \alpha^*, \lambda)\, d\pi(t) + \int_{t_h}^2 \phi_s^{-1}(t; \alpha^*, \lambda)\, d\pi(t) + \int_{t_l}^{t_h} a_p d\pi(t)
\end{aligned}$$

This computation yields:

$$a_p \int_{t_l}^{t_h} d\pi(t) = \alpha^* - \int_0^{t_l} \phi_s^{-1}(t; \alpha^*, \lambda)\, d\pi(t) - \int_{t_h}^2 \phi_s^{-1}(t; \alpha^*, \lambda)\, d\pi(t)$$

, or,

$$a_p = \left(\frac{1}{P\left(t \in [t_l, t_h]\right)}\right)\left(\alpha^* - \int_0^{t_l} \phi_s^{-1}\left(t; \alpha^*, \lambda\right) d\pi\left(t\right) - \int_{t_h}^{2} \phi_s^{-1}\left(t; \alpha^*, \lambda\right) d\pi\left(t\right)\right)$$

$$= \frac{\int_{t_l}^{t_h} \alpha^* d\pi\left(t\right) + \int_0^{t_l}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right) + \int_{t_h}^{2}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right)}{P\left(t \in [t_l, t_h]\right)}$$

$$= \alpha^* + \frac{\int_0^{t_l}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right) + \int_{t_h}^{2}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right)}{P\left(t \in [t_l, t_h]\right)}$$

■

Proof of Theorem 3.3

If $\lambda > \lambda^*$, then there exists a unique social equilibrium with incomplete separation satisfying the D1 criterion where the single central pool is located at the population average strategy, i.e., where $a_p = \alpha^*$, or equivalently, $\varepsilon\left(\alpha^*, \lambda\right) = 0$.

Proof. Theorem 3.3 follows by establishing continuity and monotonicity in α^* of the equation identifying $\varepsilon\left(\alpha^*, \lambda\right)$ and applying the Intermediate Value Theorem to identify a unique point where that equation is zero. Once this continuity is established, all that remains is to show there exists a pooling equilibrium where the pool lies below the SBP and another where the pool lies above the SBP.

$$\varepsilon\left(\alpha^*, \lambda\right) = \frac{\int_0^{t_l(a_p)}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right) + \int_{t_h(a_p)}^{2}\left(\alpha^* - \phi_s^{-1}\left(t; \alpha^*, \lambda\right)\right) d\pi\left(t\right)}{P\left(t \in \left[t_l\left(a_p\right), t_h\left(a_p\right)\right]\right)}$$

$$
\begin{aligned}
\xi\left(a_p\right) \equiv & \alpha^* - a_p \\
& = \int_0^{t_l} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) d\pi\left(t\right) + \int_{t_h}^{2} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) d\pi\left(t\right) \\
& \quad + \int_{t_l}^{t_h} a_p d\pi\left(t\right) - \int_0^2 a_p d\pi\left(t\right) \\
& = \int_0^{t_l} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) d\pi\left(t\right) + \int_{t_h}^{2} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) d\pi\left(t\right) \\
& \quad - \int_0^{t_l} a_p d\pi\left(t\right) - \int_{t_h}^2 a_p d\pi\left(t\right) \\
& = \int_0^{t_l} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) - a_p d\pi\left(t\right) + \int_{t_h}^{2} \phi_s^{-1}\left(t;\alpha^*,\lambda\right) - a_p d\pi\left(t\right)
\end{aligned}
$$

Clearly $\xi\left(a_p\right) < 0$ when $a_p < \alpha^*$ and $\xi\left(a_p\right) > 0$ when $a_p > \alpha^*$. Continuity, then, would require that $\xi\left(\alpha^*\right) = 0$. ■

Proof of Corollary 3.1

If the distribution over types is symmetric around 1, then there is a unique social equilibrium with pooling on $a_p^* = \alpha^* = 1$.

Proof. This result is proven by leveraging the symmetry of the ϕ_s inference characterizing function to show that $\varepsilon\left(\alpha^*,\lambda\right) = 0$ when $a_p^* = \alpha^* = 1$.

$$
\begin{aligned}
& \varepsilon\left(\alpha^*,\lambda\right) \\
& \quad = \frac{\int_0^{t_l(a_p)} \left(\alpha^* - \phi_s^{-1}\left(t;\alpha^*,\lambda\right)\right) d\pi\left(t\right) + \int_{t_h(a_p)}^{2} \left(\alpha^* - \phi_s^{-1}\left(t;\alpha^*,\lambda\right)\right) d\pi\left(t\right)}{P(t \in \left[t_l\left(a_p\right), t_h\left(a_p\right)\right])} \\
& \quad = \frac{\int_0^{t_l(1)} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right) + \int_{2-t_l(1)}^{2} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right)}{P(T\left(1\right))} \\
& \quad = \frac{\int_0^{t_l(1)} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right) - \int_{2}^{2-t_l(1)} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right)}{P(T\left(1\right))} \\
& \quad = \frac{\int_0^{t_l(1)} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right) - \int_0^{t_l(1)} \left(1 - \phi_s^{-1}\left(t;1,\lambda\right)\right) d\pi\left(t\right)}{P(T\left(1\right))} \\
& \quad = \frac{0}{P(T\left(1\right))} = 0
\end{aligned}
$$

■

Bibliography

GA Akerlof. A theory of social custom, of which unemployment may be one consequence. Quarterly Journal of Economics, 94(4):749–775, 1980. ISSN 0033-5533.

S. Nageeb Ali and Navin Kartik. A theory of momentum in sequential voting. Working Paper, Peer Reviewed for NAJ Economics., 2006.

LR Anderson and CA Holt. Information cascades in the laboratory. American Economic Review, 87(5):847–862, DEC 1997. ISSN 0002-8282.

J Andreoni and BD Bernheim. Social image and the 50-50 norm: A thoeretical and experimental analysis of audience effects. Econometrica, 77(5):1607–1636, September 2009.

AV Banerjee. A simple model of herd behavior. Quarterly Journal of Economics, 107(3):797–817, AUG 1992. ISSN 0033-5533.

JS Banks and J Sobel. Equilibrium selection in signaling games. Econometrica, 55(3):647–661, MAY 1987. ISSN 0012-9682.

M Battaglini. Sequential voting with abstention. Games and Economic Behavior, 51(2):445–463, MAY 2005. ISSN 0899-8256. doi: 10.1016/j.geb.2004.06.007.

BD Bernheim. A theory of conformity. Journal of Political Economy, 102(5): 841–877, OCT 1994. ISSN 0022-3808.

S Bikhchandani and S Sharma. Herd behavior in financial markets. IMF Staff Papers, 47(3):279–310, 2001. ISSN 1020-7635.

S Bikhchandani, D Hirschleifer, and I Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy, 100(5):992–1026, OCT 1992. ISSN 0022-3808.

IK Cho and D Kreps. Signaling games and stable equilibria. Quarterly Journal of Economics, 102(2):179–221, MAY 1987. ISSN 0033-5533.

IK Cho and J Sobel. Strategic stability and uniqueness in signaling games. Journal of Economic Theory, 50(2):381–413, APR 1990. ISSN 0022-0531.

D Chowdhury and D Stauffer. A generalized spin model of financial markets. European Physical Journal B, 8(3):477–482, APR 1999. ISSN 1434-6028.

HL Cole, GJ Mailath, and A Postlewaite. Social norma, savings behavior, and growth. Journal of Political Economy, 100(6):1092–1125, DEC 1992. ISSN 0022-3808.

HL Cole, GJ Mailath, and A Postlewaite. Investment and concern for relative position. Review of Economic Design, 6(2):241–261, September 2001. ISSN 0022-3808.

R Cont and JP Bouchaud. Herd behavior and aggregate fluctuations in financial markets. Macroeconomic Dynamics, 4(2):170–196, JUN 2000. ISSN 1365-1005.

E Dekel and M Piccione. Sequential voting procedures in symmetric binary elections. Journal of Political Economy, 108(1):34–55, FEB 2000. ISSN 0022-3808.

A Devenow and I Welch. Rational herding in financial economics. European Economic Review, 40(3-5):603–615, APR 1996. ISSN 0014-2921.

D Fudenberg and J Tirole. Game Theory. The MIT Press, Cambridge, MA, 1991.

David A Hirshleifer and Siew Hong Hong Teoh. Herd behavior and cascading in capital markets: A review and synthesis. Dice Center Working Paper No. 2001-20, (20):Available at SSRN: http://ssrn.com/abstract=296081 or DOI: 10.2139/ssrn.296081, 2001.

Angela Hung and Jeff Dominitz. Homogeneous actions and hetergeneous beliefs: Experimental evidence on the formation of information cascades. Econometric Society 2004 North American Winter Meetings, 64, 2004.

E Kohlberg and JF Mertens. On the strategic stability of equilibria. Econometrica, 54(5):1003–1037, SEP 1986. ISSN 0012-9682.

J Lintner. The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets. Review of Economics and Statistics, 47(1): 13–37, 1965. ISSN 0034-6535.

A Postlewaite. The social basis of interdependent preferences. European Economic Review, 42(3-5):779–800, MAY 1998. ISSN 0014-2921.

J Rohlfs. Theory of interdependent demand for a communications service. Bell Journal of Economics, 5(1):16–37, 1974. ISSN 0361-915X.

WF Sharpe. Capital Asset Prices - A theory of market equilibrium under conditions of risk. Journal of Finance, 19(3):425–442, 1964. ISSN 0022-1082.